Project Final Report
Reward Structures

Archie Chapman (contact author), Sleiman Mhanna,
Gregor Verbič
The University of Sydney

Archie.Chapman@sydney.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.

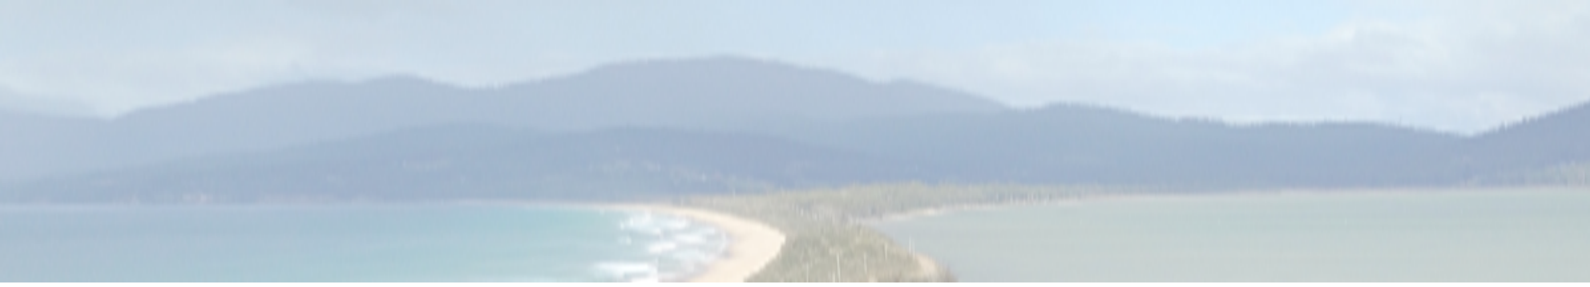

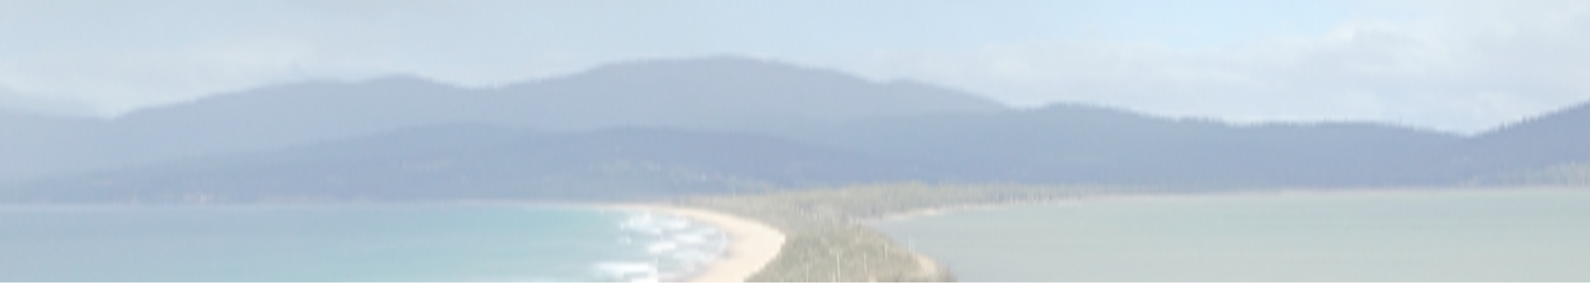

Executive summary

Reward structures refer to the design of payments to customers for the network support services that they provide. In the CONSORT project, we have developed reward structures that unpack the “value-stack” available to distributed energy resources. In keeping with recent microeconomic reforms to the electricity sector, these reward structures implement value-reflective pricing methods for network support, so align with the move to cost-reflective network tariffs put in place after the AEMC’s *Power of Choice* review.

A solution to the problem of pricing the network support provided by customer-owned batteries was found in the economic concept of the *Shapley value*, which was used as a template of an ideal reward structure. The Shapley value provides a principled set of properties related to network support pricing, most importantly a form of fairness (equal treatment to equal contributions and independent pricing of independent effects) and efficiency (full disbursal of the rewards available). However, since directly using the Shapley value reward structure in practice is computationally infeasible, the project developed various estimation and approximation methods. These were integrated with the NAC algorithm, and successfully deployed in the field. Analysis of the payments computed by these reward structure methods indicated that they did indeed reflect the batteries’ value to the network in principled ways, with useful findings for distribution network companies and retailers.

Despite these successes, the reward structure methods developed had varying degrees of success by practical computational metrics. One finding from the reward structures work package is that the exceptionally difficult task of calculating the Shapley value of a network support event makes it infeasible to use as a method of generating spot or even close-to real-time prices, unless severe approximations of the computation are made. Additionally, although they could be deployed in the CONSORT trials, the required approximations undermine the use of these reward structure for calculating customer payments in more complicated problems of sharing multiple DER value streams, for example, when simultaneously managing network voltages as well as thermal limits.

Nonetheless, a path forward to the use of value-reflective reward structures in paying for network and power system support services has been plotted based on the findings of the project. The methods developed can overlay any DER control scheme, regardless of its level of sophistication. We will continue to develop the required models and methods, and to prosecute the arguments, for value-based reward structures for support provided by customer-owned DER.

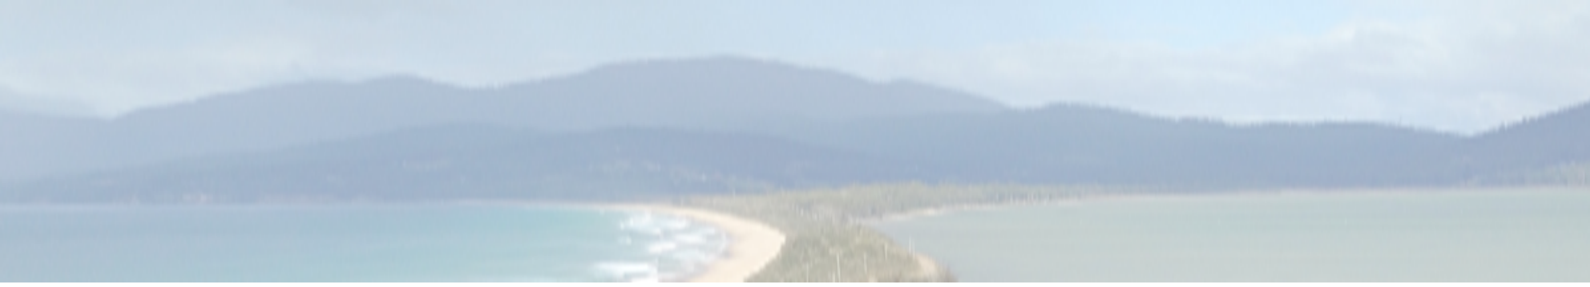

Table of contents

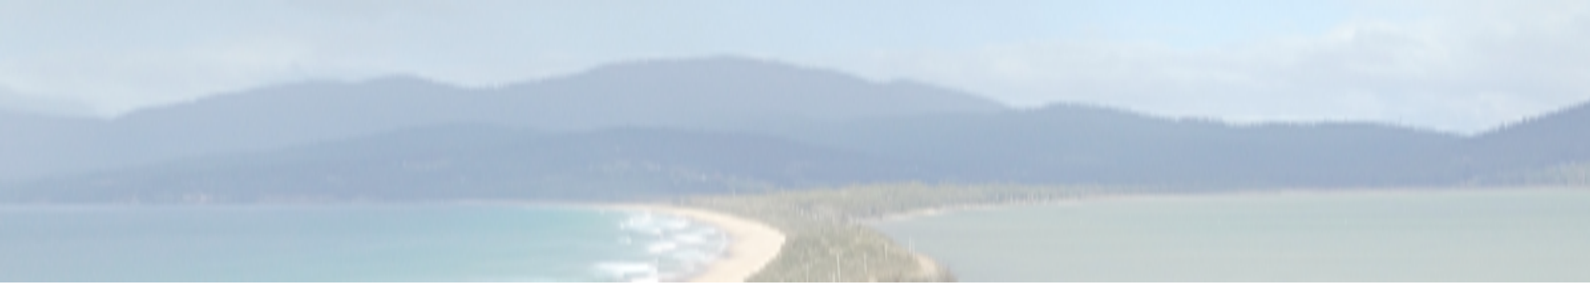

1 Introduction

This report considers the design payments to the customers involved in the Bruny Island battery trial, for the network support provided by their batteries. We call these payment designs *reward structures*. We have designed the reward structures to be value-reflective, in keeping with the move to cost-reflective network tariff structures that have been put in place in response to the AEMC's *Power of Choice* review.

Specifically, the reward structures work package investigated methods for non-linear pricing of network support provided by any behind-the-meter energy resources, including but not limited to residential batteries. These methods are based on cooperative game solution concepts, and as part of the project, the Shapley value was selected as the best suited method of generating value-reflective prices for network support.

The report is structured as follows. The next two subsections provide the rationale for adopting a cooperative game pricing problem formulation and selecting the Shapley value as a solution. Following this, Section 2 considers the computational methods developed during the project used to implement the reward payments. Section 3. Finally, we provide some concluding thoughts and recommendations for further research in this area.

1.1 Cooperative games

Cooperative games model surplus division problems, in which a group of *players* or *agents* that have agreed to cooperate to earn a joint reward, also have to decide how that reward is allocated among them. That is, in cooperative game theory a *coalition* should divide its reward among its members. Specifically, these games are used to model strategic situations involving rational, self-interested agents that can form binding contracts with one another to pursue a common action. Thus, the problem one faces in analysing a cooperative game is to find a division of the rewards earned by a coalition that, first, ensures stability of the coalition, and second, achieves some distributional goals, such as fairness, proportionality, etc. It is the ability to commit to a course of action that distinguishes cooperative games from the more-widely known *non-cooperative* game formulation (e.g. which are used to model auctions or congestion problems).

Pricing rules based on cooperative game models have some key differences with conventional pricing methods employed in auctions or used in charging for the use of regulated network assets. For example, marginal-cost pricing is the typical approach used to price power generation. In an optimisation framework, marginal-costs are derived as dual variables or constraint multipliers. In the NAC, this could give rise to locational marginal prices that are computed by considering the incremental benefits of discharging customers' batteries on diesel costs. In this, storage assets are assumed to be available all the time, and pricing is

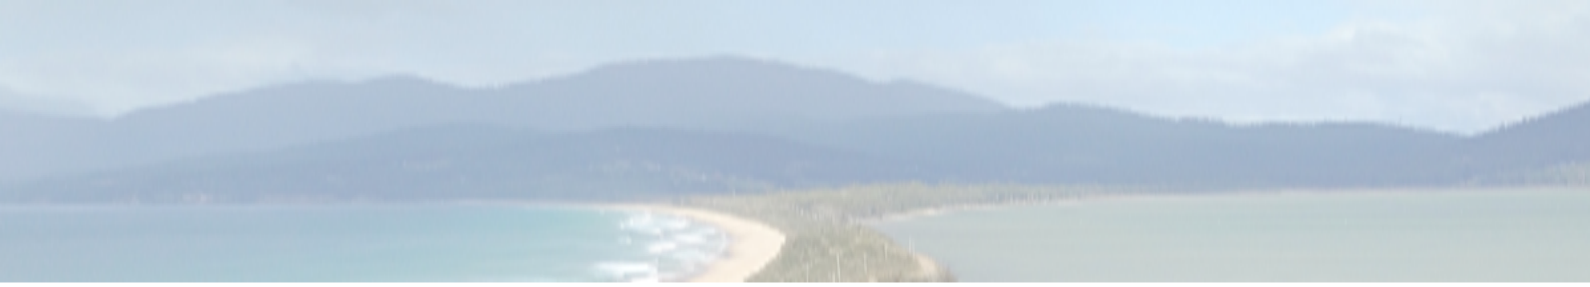

spot pricing (i.e. in five minute intervals). Effectively, prices are an artifact of the optimisation model formulation used. In contrast to the incremental benefit pricing approach, cooperative solutions are derived based on certain properties or axioms, which can be chosen to encode principles, such as fairness and/or efficiency. This allows for some greater flexibility in the way prices can be derived in a given setting, including allowing for redundancy to be priced, or entire dispatch schedules to be allocated a price representing their benefit as a whole, rather than as the sum of the prices in the spot-market. Of particular interest to us is the principled way that cooperative games allow us to isolate the self-interested behaviour of customers from their NAC-coordinated network support actions. Our results later illustrate how the Shapley value provides us with a way to solve these economic problems.

Formally, the cooperative game model and its Shapley value solution can be straightforwardly mapped to the network support setting, as follows:

- A coalition of battery controllers (the agents) is paid by a DNSP to provide enough load relief to overcome a predicted thermal constraint excursion. Note that the actual 'players' in this game are the battery controllers acting as agents for householders or customers, who rely on the technology of Reposit and the NAC to negotiate and make decisions on their behalf and in their best interests.
- In practice, the battery owners agree to perform some optimal joint network support action using a form of distributed optimization and control platform, and in the analysis we present in this report, we draw on NAC algorithms developed by ANU (see Network Aware Coordination FInal Report [1]).
- For completing this network support action, the aggregated coalition of battery systems as a whole receives a reward in the form of a payment from TasNetworks. In the Bruny Island Battery Trial (the Trial), the size of this payment is determined using a heuristic. However, in the future we reasonably expect it to be determined by a network monopoly regulating body, such as the Australian Energy Regulator.
- At the same time, each DER owner incurs some private cost, in the form of energy cost savings foregone, round-trip losses due to charge and discharge inefficiencies, and device degradation.
- Thus, the players in the cooperative game are DER owners, and the payment to the coalition has to be divided among the DER owners.

For this payment division problem, we investigated the use of the Shapley value solution.

Before introducing the Shapley value, some definitions are required. If the players in a cooperative game agree to work together, they form a *coalition*. If all *N* player form a coalition, it is called the *grand coalition*. Collectively, the joint action has some *worth* associated with it (i.e. revenue). The characteristic function of a game defines the worths of all possible coalitions, $S \subseteq N$, and is denoted $w(S):\ 2^N \to \Re$ with the worth of the empty coalition equal to zero, $w(\emptyset) = 0$.

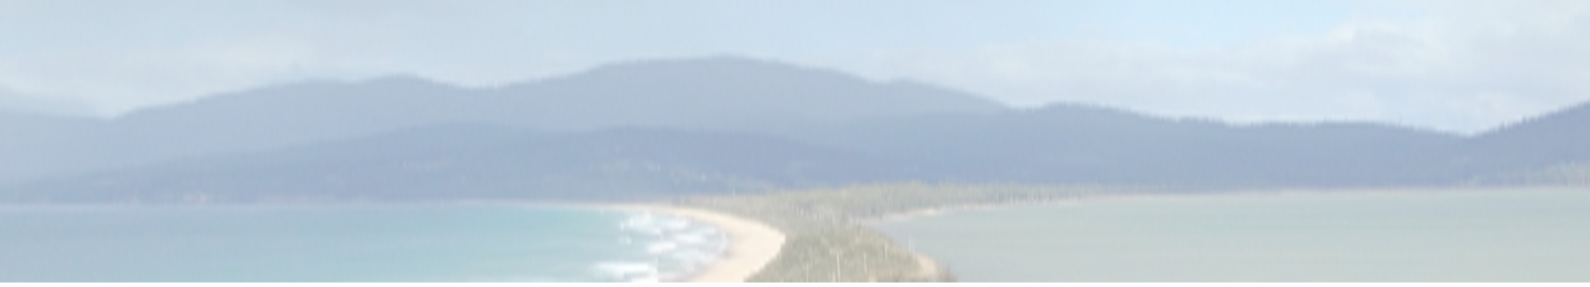

1.2 The Shapley value

The Shapley value is a seminal cooperative value division rule developed by game theory researchers. It was derived to satisfy the following properties (referred to as “axioms” in the research community):

- Efficiency: the full payment is allocated;
- Symmetry: identical DER are allocated the same amount;
- Additivity: if an agent is involved in two separable games at once, its allocation is the sum of its payoffs in the two games separately, and;
- Null player: those who contribute nothing to the coalition receive zero payoff.

These properties make the Shapley value an ideal metric against which to evaluate the value- (or cost-) reflectivity of a particular “simple” reward (or tariff) structure. In addition, and by design, the Shapley value is unique and always exists, which is uncommon for cooperative solution concepts.

The Shapley value for player/customer *i* is given by the function:

$$\varphi_i(w) = \frac{1}{n} \sum_{S \subseteq N \setminus i} \binom{n-1}{|S|}^{-1} \left(w(S \cup i) - w(S) \right)$$

Here, the value function φ_i has the following intuitive interpretation: consider a coalition being formed by adding one player at a time. When *i* joins the coalition *S*, its *marginal contribution* to the resulting coalition is given by $w(S \cup i) - w(S)$. This is the last part of the expression above. Then, for each player, its Shapley value payoff is the average of its marginal contributions over the possible different orders (or permutations) in which the coalition can be formed, where the binomial term is the number of coalitions of size $|S|$.

The main challenge in computing the Shapley value is that the number of possible coalitions grows exponentially with the number of players (i.e. 2^N). In the CONSORT project context, this means the *exact computation* of the Shapley value requires solving 2^N or more than 34 billion optimal power flow problems. Clearly this is infeasible, so a major focus of the rewards structure work package has been to investigate ways to efficiently approximate the Shapley value. The next section describes two approaches developed during the project to overcome this computational hurdle.

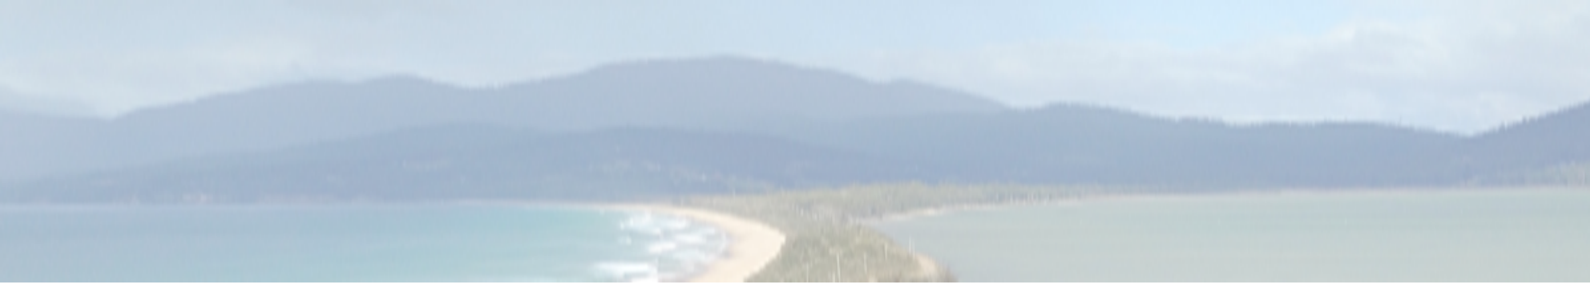

2 Reward structure computation algorithms

The first part of the reward structures work package was to develop algorithms to compute value-reflective rewards for customers contributing to the NAC network support. As discussed above, the approach adopted was to use the Shapley value as reward structure template, and the supporting computational procedures for this are described here. This proved to be a very technically-challenging task, and several approaches were trialed over the course of the project. The two most-effective are documented in the section.

2.1 General Shapley value-based approach

The steps involved to develop the general Shapley value-based reward structures approach were as follows:

1. Compute the Shapley values for each customer (using either a sample-based or relaxation approximation method) in terms of diesel generator cost reductions;
2. Normalise the Shapley values, so that they represent the proportion of the TasNetwork budget to allocate to each customer;
3. Calculate the TasNetworks total expenditure, given as the total battery kWh discharge during the peak period multiplied by $1/kWh, and allocate this amount to customers in proportion to their normalised Shapley values, and
4. Apply a lower bound to ensure no customer was worse off for providing support.

Two algorithms for Steps 1 and 2 regarding the Shapley values are discussed in detail in Sections 2.2 and 2.3 below.

For step 3, the TasNetworks budget was set by calculating the sum of battery discharging during the support period T_{sup} times $1/kWh:

$$TN\ Budget\ =\ \$1.0\ \times \sum_{i \in N} \ \sum_{t \in T_{sup}} P_{i,d}(t)$$

where $P_{i,d}(t)$ is the discharged energy from battery $i \in N$ during dispatch interval $t \in T_{sup}$.

For step 4, the lower bound on payments over the period is given as the amount which will ensure that, no matter what the battery does during the actual peak, the customer will not be worse off. Typically this is $1.81 per battery discharge event for the LG Chem batteries widely in use in the trial (derivation provided in Section 3.1).

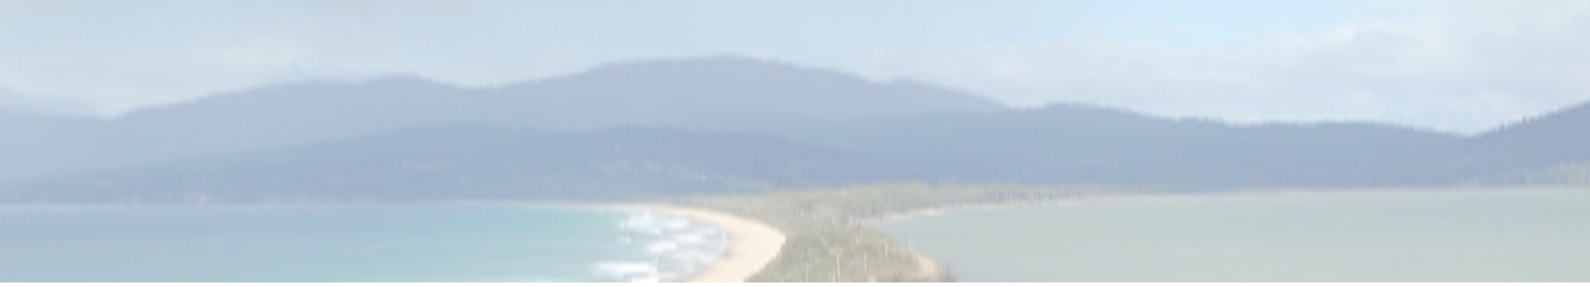

2.2 Algorithm 1: Sample-based estimation of the Shapley value

The sample-based Shapley value computational architecture used to compute the Shapley value of customers' contributions under the NAC is found in Figure 1 below.

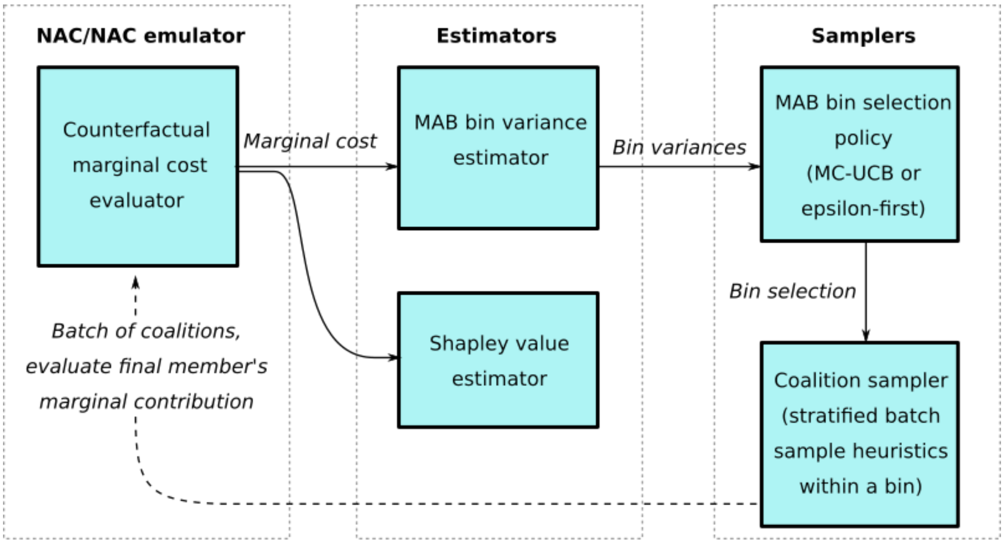

Figure 1. The computational architecture used to compute the Shapley value of customers' contributions under the NAC

NAC emulator: In this part, an optimal power flow solver computes the total cost of the NAC (diesel and customer retail costs) with and without a particular customer's participation. The difference in cost give the customer's *marginal contribution* to the NAC assuming only a sub-coalition $S \subseteq N$ of customers participate. This is the fundamental measurement used to estimate the Shapley value. For this step, accelerated NAC-emulating optimisation algorithms have been developed. More on this work can be found in Section 2.4.

Estimators: The marginal contributions are used to drive two different estimation routines. The first is a *multi-armed bandit* (MAB), which is used to sample sub-coalitions and estimate the Shapley values, and the Shapley value estimator itself.

In more detail, the marginal contributions are used to drive the online sampling algorithm, which uses the MAB framework for online (or iterative) optimisation of unknown stochastic functions (upper arm of the diagram above). MAB bins are constructed for groups of the sub-coalitions *S* arranged by size, and the variance of the marginal contributions for each bin size is recorded at each iteration. The marginal cost that is computed by the NAC emulator is used to update a running estimate of the marginal cost variance for its respective MAB bin. The bin variances are used by an MAB policy to select the next bin (sub-coalition size) to sample from and evaluate. At the same time, the marginal contributions are used to update a vector of values of the Shapley value estimate, one for each customer. These are standard running average and variance estimates.

Samplers: A range of MAB bin selection heuristics have been evaluated in the course of the project. We used a well-known and widely-used method, called the *upper confidence bound*. In the setting of variance-based efficient estimation, this policy selects the bin with the largest variance; which in our setting is the coalition size with the largest variance on the average marginal contribution.

Given a coalition size, the coalition sampler is then used to sample a particular batch of coalitions of that size. The specific mechanism used begins by randomly selecting a *pivotal* sub-coalitional of the required size, to evaluate as a base case for all marginal coalitions. Then a full batch of sub-coalitions is generated by, first, removing members of the pivotal sub-coalition one at a time (so that their marginal contributions can be computed) and second, adding in each customer not in the pivotal coalition. This produces *N+1* sub-coalition problems for solving by the NAC emulator and is used to return *N* marginal contribution samples, one for each customer.

Putting these modules together, the overall process is as follows:

1. The MAB keeps track of the variance of marginal contributions to coalitions of each size, and the MAB bin selection routine choose the coalition size with the greatest variance on each sampling iteration;
2. After selecting a coalition size, the coalition sampler routine uses a stratified batch sampling heuristics to select a number of coalitions with the corresponding coalition range;
3. The batch of coalitions is passed to the NAC counterfactual marginal cost evaluator, which returns a batch of marginal contributions, one for each customer;
4. The marginal contributions are passed to (i) the MAB routine, which returns to step 1, and (ii) the Shapley value estimator, which updates its estimates;
5. The process is repeated until Shapley value estimates reach the required degree of accuracy.

2.2.1 Discussion on computational performance

During the trial, the Shapley values were to be computed using the sample-based estimation method as described above. This guarantees convergence to the true Shapley values *in the limit*, that is, as the number of sample grows large. Nonetheless, for 34 customers, this takes a very long time to compute exactly, so the estimation is truncated early when sufficient statistical confidence is attained. Usually this is up to 200 hours of computation, corresponding to around 100-150 samples for each of the 34 customers (i.e. ~5000 NAC instances), and only in certain unusual circumstances is it reliably quicker. Even when making use of high-performance algorithms and multi-core computing hardware (as discussed in Section 2.4), this can only be brought down to around 10 hours of computation.

As such, the sample-based approach took too long for the trial's messaging timelines, which had a computational wall time duration of around 1-2 hours. This was a major learning during the project, as the limits of the sample-based Shapley value computation methods were not well understood before we began the CONSORT trial. The long computational times of the sample-based approach also drove us to explore various approximations of the Shapley value. One approach - an online heuristic - stood out for meeting project requirements, and was ultimately integrated into the NAC computation stack and deployed during the trial, as described in section 2.3.

2.3 Algorithm 2: Online heuristic Shapley value approximation

The online heuristic Shapley value approximation works by (i) applying a mathematical linearisation or "averaging" to the contribution of the batteries on each phase to the network support problem, which is effectively linearising the network losses component, and (ii) aggregating phases in the counterfactual computation to reduce the size of the Shapley value computation of the problem by several orders of magnitude. Taken together, this produces an algorithm that runs in seconds. These two components are described in detail below.

Linearisation: The purpose of the linearisation is to determine an analytically tractable marginal contribution of a customer to the network support problem. The marginal contribution depends on the customer's available energy and power, and its voltage ratio, whether it is connected to the line that is setting diesel dispatch level (if it is dispatching), and the amount of diesel reduction that is available to the NAC system.

The linearisation has three steps. First, the network loss effects on the network support problem are lineralised around the NAC operating point, as follows. Given the active power component at a generator, and current down a line, of a NAC solution, *P** and *I**, respectively, we derive the following relationship. If we fix the voltages of customer *i*'s bus (bus *A*), V_A, and the diesel generator, V_g, at the NAC operating point, we can express the power flow, including an approximation of the losses to the network, in terms of the currents, as a linear function of the losses. With some manipulation of terms, and assuming the change in current from a battery is full compensated by a change in the generator current $I^*_g - I_g \approx I^*_i - I_i$, we have:

$$P^*_g - P_g \approx (I^*_i - I_i)V_g \approx (P^*_i - P_i)V_g/V_A$$

Denote the value V_g/V_A the *voltage ratio* for *i*. This value is used to account for network losses in the network support cooperative game (including all auxiliary games to be defined in what follows).

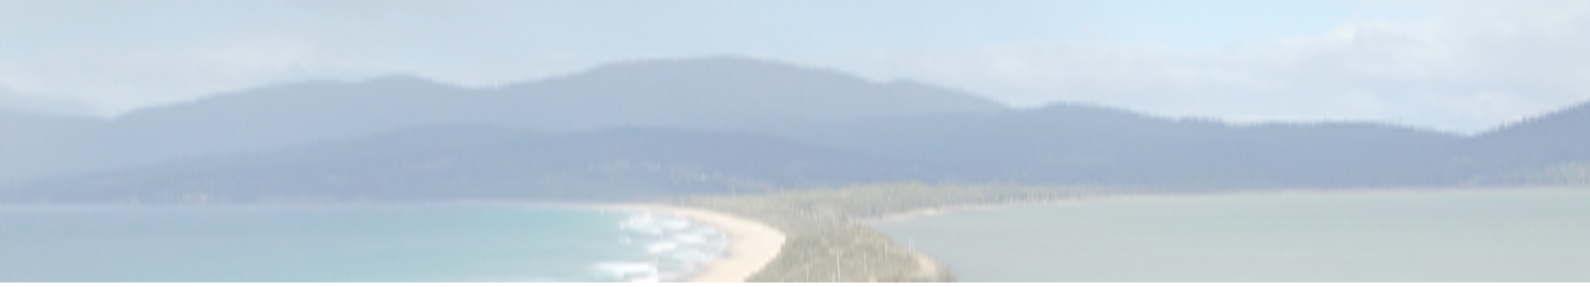

The voltage ratios are fixed for the cooperative game, and are computed immediately after the peak event based on measurements of the actual NAC operation.

Second, a *network_overlimit* value is computed, to represent the potential diesel use reductions available to the NAC. This value is computed from a simulation of the network effect of removing the diesel generator and not operating the NAC. This results in line currents that exceed their limits, and gives a counterfactual case for evaluating the benefit of the diesel generator and the batteries. Specifically, the *network_overlimit* value is the sum of the current flowing on the highest-current line during the peak time over the current limit in the counterfactual simulation, converted to an energy value (i.e. convolved with the voltage profile during the simulated peak). For each time-slot, this value is trimmed to the maximum collective power of the batteries whenever the required power exceeds the power capacity of the batteries. (A potential further extension is to also trim to the value to account for the diesel minimum dispatch level, although this step has not been implemented.)

Third, to generate this marginal value for each customer, we define two values on the customer side:

- *B_E_nac* is the total energy available (to the NAC) from each battery, and
- *B_E_selfish* is the energy used under a selfish schedule (i.e. no NAC).

Then the fixed marginal contribution of customer *i*, denoted θ_i, is:

$$\theta_i = (B\_E\_nac - B\_E\_selfish)\, V_g / V_A$$

where A is the bus connecting i to the network. This is an approximation of the additional energy available to the NAC at the generator bus over what would be selfishly dispatched by the batteries during the peak period; or effectively the extra energy that the NAC pulls out of the batteries to offset the diesel generator (or that the NAC can dispatch in any particular counterfactual coalition instance). The θ_i value is used to apportion the *network_overlimit* value between the NAC-enabled customers.

Aggregation: We can treat θ_i as customer *i*'s linear component of the game's characteristic function. Once the cooperative game is a linear game, we can use some analytic tricks to efficiently compute the SV. In particular, we consider block coalitions of customers grouped by phase sans *i*, and then compute *i*'s individual contribution to these block coalitions to produce a final approximation of the SV.

In more detail, the aggregation step exploits the linearity property of the SV to massively reduce the number of marginal value calculations required. For each player, an auxiliary game is constructed by aggregating the remaining customers' network support by phase, and treating each aggregated phase as one player in the auxiliary game. Each auxiliary game has four players, so its characteristic function has 16 entries, and therefore the SV of the focal

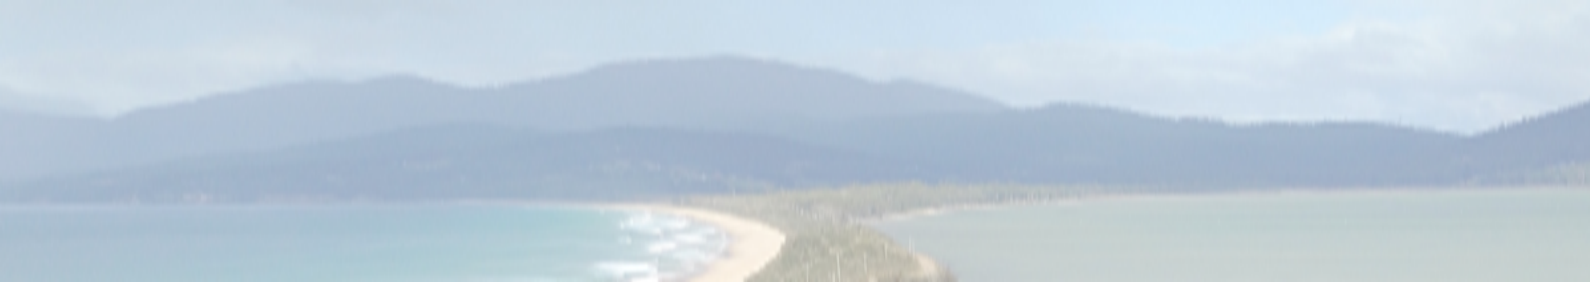

player in this auxiliary game can be computed exactly with no difficulty and is solved algebraically.

One potential source of error in this aggregation step is that the “merit order” effect induced by customers having different loss reductions, due to their location on the network and the distribution of loads, is lost in the by-phase averaging. A second is that the phases with many customers may be receiving a disproportionately large SV, due to the relatively large size of their aggregated theta’s. The comparison of this heuristic to the sample-based method mis conducted to assuage these concerns.

2.4 High-performance DER coordination algorithms

Realistic and fair reward structures require the modelling of the Bruny Island distribution network and subsuming it in the Shapley value (SV) approximation scheme. A detailed model of the network is important because the dispatch schedule from the batteries is directly affected by *voltage drop*, *congestion* on the undersea cable, and the *phases* that this battery is connected to. The resulting framework is a multiperiod optimal power flow with flexible demands, otherwise known as *Network ware Coordination (NAC)*. The NAC is a large-scale nonlinear nonconvex programming problem that has an NP-hard computational complexity. In fact, for a 12-hour scheduling horizon and a half-hourly time resolution, the NAC problem has more than 200,000 variables and more than 230,000 constraints. The SV approximation requires pivoting in and out a large number of different coalitions of agents. Therefore, it is necessary that the NAC solved at each instance of this pivoting is fast enough to ensure computing the reward structures within a feasible time frame. The NAC can be solved in two ways, either centrally or in a distributed fashion.

The centralized NAC is coded in *Python* with the use of *MADOPT* as an automatic differentiation tool for the backend solver *IPOPT*. In the aim of speeding up the solve time of the NAC, line current limit constraints are only considered for the undersea cable, which greatly reduces the number of constraints as opposed to enforcing line current limit constraints on *every* overhead line in the network. This stems from the observation that most of the lines in the network have current limits greater than that of the undersea cable. This means that the undersea cable will always hit its thermal limit before the other overhead lines downstream from it in the network do. Consequently, this reduction in the number of quadratic constraints results in a speedup of up to 5 times. Specifically, before the constraint reduction the centralized computation time was about 300 seconds on average compared to 60 seconds on average after the constraint reduction.

Another way of solving the NAC problem is to duplicate the complicating variables which allows for a component-based decomposition of this problem in the dual domain. The resulting dual problem can now be solved using the *alternating direction method of multipliers* (ADMM). The advantages of ADMM are 1) that it is robust and; 2) that it preserves the separable structure of the problem. As a result, the problem can be solved in a distributed fashion, which has the potential of being faster than the centralized approach. Unfortunately, achieving this superior performance requires a large number of cores to be used for the parallel computation.

We emphasise that the computational complexity of the problem manifests as a barrier to Shapley value computation because of the extremely large number of NAC instances that need to be solved. In contrast, in the operational version of the NAC, described in [1], only a single NAC instance is solved for each time step, and the methods developed by ANU overcame the computational challenges this presented.

3 Reward structure trials

The CONSORT reward structure trials considered two independent design considerations for the network support payments made to customers. One major consideration was the the design of payments made to customers. This includes the timing of peak event alerts and payment information sent to customers. In order to understand these effects two treatments were developed and trialed: 1) *Energy Reserve* payment and: 2) *Energy Use* payment. Roughly speaking, *Energy Reserve* payments were computed on forecast NAC operation and communicated to Trial participants via a text/email notification received via the Reposit system before a forecasted peak event; while *Energy Use* payment were computed from actual NAC operation data and communicated to participants after a peak event. Section 3.2 discusses the findings from these payment design trials.

The focus of Section 3 of this report was the technical development of computational routines used to implement Shapley value-based payments. In particular, Algorithm 2, described earlier, was integrated with the live NAC deployment as part of the field trials; this work is discussed in Section 3.3.

Finally, in Section 3.4, we analyse the network support payments and financial outcomes for the customers.

To begin, however, we provide an overview of the trial periods, during which customers were paid for their network support.

3.1 Trial events

The reward structures were trialed alongside the NAC during the peak periods listed in Table 1 below (a subset of the trials reported in the Network Aware Coordination Final Report [1]).

Dates	Details	Result
29/03/2018 to 03/04/2018	Easter holiday. *Energy Reserve* payment type with heuristic reward in use.	6 peaks. *Energy Reserve* payments and customer messaging executed without problems.
13/04/2018 to 03/05/2018	April school holidays.	0 peaks, so no payments.
08/06/2018 to 12/06/2018	Queen's Birthday long weekend.	5 peaks.
13/07/2018 to 23/07/2018	July school holidays. Switched to *Energy Use* payment type	5 peaks. *Energy Use* payment and customer messaging executed

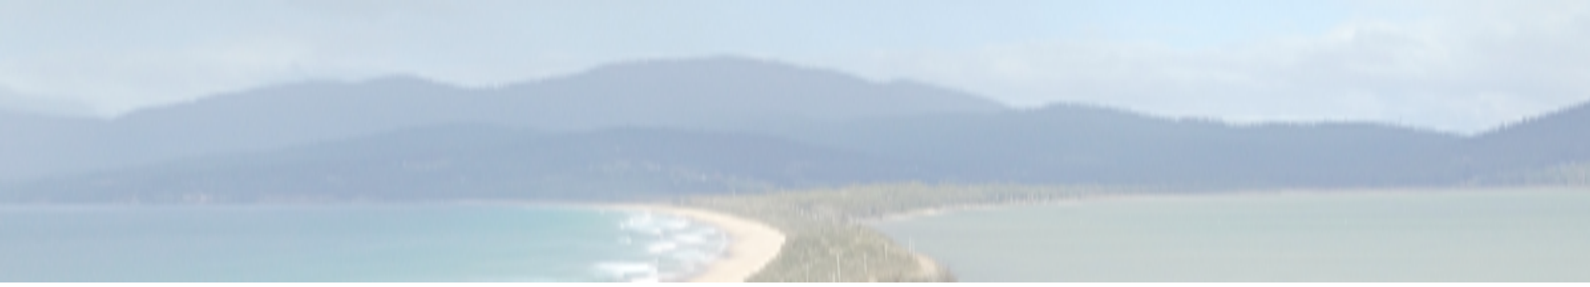

	(with heuristic reward) on 15/7/18.	without problems from morning of 15 July.
19/12/2018 to 09/01/2019	Summer holiday period. Continued with *Energy Use* payment type.	0 peaks, so no payments.
08/02/2019 to 09/02/2019	Special session to live test Shapley-value based reward structure code.	Successful demonstration of SV RS code.

Table 1. A summary of the trial periods where the team tested the reward structures.

Some general features of all of the trials are described next.

Heuristic payment: Due to delays in developing Shapley value routines for the live trials, we used a heuristic payment for the trial. This payment was based on battery discharging during the support period T_{sup}, and is given by:

$$Provisional\ payment\ =\ \$1.0\ \times \sum_{t \in T_{sup}} P_d(t)\ .$$

where $P_d(t)$ is the discharged energy during dispatch interval $t \in T_{sup}$. In Section 3.2, the difference between the *Energy Reserve* and *Energy Use* is explained in terms of whether the power discharged value is a forecast or an actual value.

Payment lower bound: As noted previously, a lower bound on payments was always applied, and the bound applied in the two test cases described here is the same. The lower bound is set to an amount that over the period, will ensure no matter what the battery does during the actual peak, the customer will not be worse off, and is given by:

$$Payment\ bound\ = \$0.343\ \times\ min[SoC_{max}, |T_{sup}| \times P_d^{max}\ ]$$

where:

- $0.343 is the tariff difference for buying high and selling low over the support period divided by the inverter round-trip efficiency, 0.9),
- SoC_{max} is the maximum energy that can be stored in the battery, and
- $|T_{sup}| \times P_d^{max}$ is the most that can be discharged during an interval of length $|T_{sup}|$.

In other words, this bound is the approximate difference in retail electricity bills between the best and worst case uses of the battery from the customers' perspective. This was the same bound as used previously.

Budget: The TasNetworks budget was set by calculating the sum of battery discharging from all customers during the support period T_{sup} times $1/kWh:

$$TN\ Budget\ =\ \$1.0\ \times \sum_{i \in N} \sum_{t \in T_{sup}} P_{i,d}(t)$$

where $P_{i,d}(t)$ is the discharged energy from battery $i \in N$ during dispatch interval $t \in T_{sup}$.

3.2 Payment types: *Energy Reserve* versus *Energy Use*

Under the *Energy Reserve* payment, CONSORT paid customers based on how much of a customer's battery capacity we expected to use, advised in advance of the peak event. The *Energy Use* payment paid customers based on how much energy was actually used from their battery to provide the network with support, advised after the event.

Payment type trial supporting information: In order to understand Trial participants' perspectives on the two treatments, payment timing coincided with social science and NAC activities. In particular they were timed with the data collection via 'energy diaries' and interview processes undertaken by CONSORT's social science team. The reward structures team also assisted in preparing communications conveying this information for customers, including emails, FAQs, website blog posts and video explainers. We used the following text to explain the difference between the two payment types to participants:

> *"For the **Energy Reserve** payment type, we will be paying you before an event to reserve some energy in your battery. Think of it like we are borrowing your battery for a little while. We pay you based on how much energy we need you to reserve in your battery, where you are on the electricity network, and what sort of system you have. Importantly, we will tell you how much we will pay you before the event and pay you regardless of whether we actually use your battery or not. Thinking again of the car analogy, this time we pay you to reserve your car for our use for a bit, regardless of how much we end up using it, and we pay a different amount for each car depending upon how useful we think it will be to us.*
>
> *" [Under the] **Energy Use** payment… We will tell you how much you were paid after an event and it will be based on how much energy we buy out of your battery. Similar to the Energy Reserve payment system though we will vary how much we pay you based on 'usefulness', what sort of system you have and where you are in the network. Again, if you think of it as us borrowing your car, you'll be paid based on how far we actually needed to drive your car, with each car owner being paid different amounts."*

3.2.1 Findings

From the participants' point of view, receiving notifications of the size of network support payments before or after the events - that is, as energy reserve or energy use payments, respectively - appeared to have very little salience or effect on behaviour (this finding is discussed in more detail in the Social Science final report [2]). This is useful knowledge, because we had previously conjectured that some customers may not like to have their battery used for network support without knowing the payment they would receive in advance. Indifference between the two indicates that customers may be happy to be paid in any reasonable form, or it may be an indication that other factors, such as having trusting relationship between the parties or a deeper level of understanding of the technology, may have greater salience in customers' minds.

From the DNSP and aggregator point of view, however, the two payment types present quite a different challenges, the main being the veracity of forecast information. Specifically, under the *Energy Reserve* payment type, the payment amount depends on a forecast time series of each customer's energy use, and of the network state, around the forecast peak period. This brings with it two main difficulties.

First, forecasting peaks is in itself a challenge, as they are rare events, and the NAC settings are biased to prepare for a peak even if one does not eventuate. This lead the *Energy Reserve* payment to allocate payments for peak events that do not realise. In contrast, *Energy use* payments do not suffer from this "false peak" problem, as they are computed after the fact.

Second, using customer load and network state forecasts to run optimal power flow studies on which payments are based is invariably inaccurate, as the power flow solutions are notoriously sensitive to network conditions. At the same time, load forecasts for individual households are very volatile, as there is no large number effect smoothing the load (e.g. compared to load profiles used for network planning problems, where the large number of customers smooths out their individual volatility). Within this context, we noted that the network voltage effects on the battery merit order are of importance, which can vary dramatically depending on load values, and as a consequence can produce *Energy Reserve* payments that does not closely correspond to the actual energy discharged from batteries for network support.

Beyond issues of forecasting and accuracy, a third challenge regards the potential for strategic behaviour by customers under the *Energy Reserve* payments (and to a much lesser extent the *Energy Use* payments with the heuristic payment, but not the Shapley value rewards). However there is a commitment problem between when the customer is paid and when they use energy in their homes, that potential allows customers to be paid for reserving energy, but does not bind them to keeping that energy available for use during the upcoming network support period. Instead, the energy could be used to serve local load, thereby removing the benefit to the network. We stress that this was not observed in the trial, but it was a concern uncovered during the trial by the CONSORT team, and could be used by sophisticated customers with significant or complete agency over their loads that are involved in NAC-like network support or demand response operations.

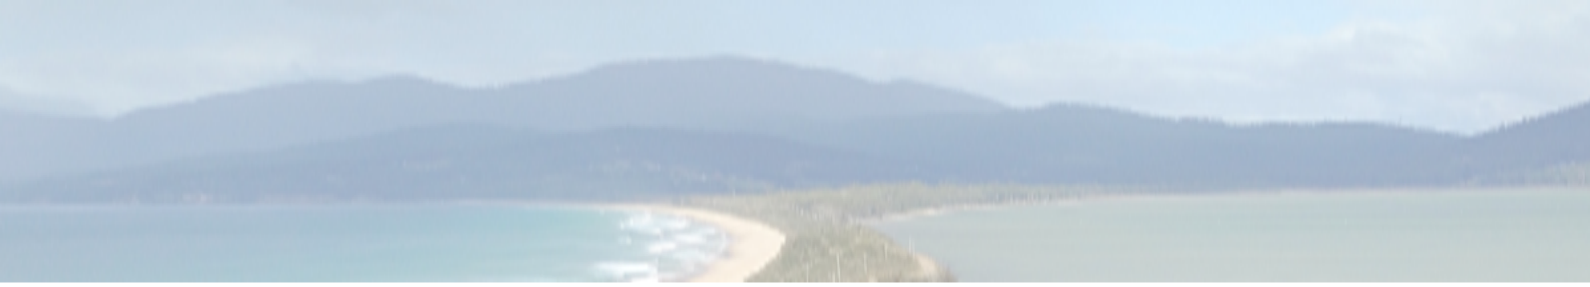

These three implementation difficulties, and the aforementioned lack of overt preference from customers, makes the case for recommending *Energy Use* payment type over *Energy Reward*, although there may be other factors of context in which the exchange is taking place that could void this finding.

3.3 Online Shapley value payments: Integration of Algorithm 2 with NAC deployment

The analytic Shapley value heuristic (Algorithm 2, described in Section 2.3) went through several iterations of refinement before being deployed by ANU in the NAC software stack. The reward calculation procedure was triggered to run immediately after any peak event. This involved the following steps:

1. The required data was collected from the actual NAC run, including customer PV generation, load profiles, and initial battery state-of-charge data, background loads at MV feeder buses, diesel dispatch information, bus voltages and line currents;

2. The network effects of removing the diesel generator and not operating the NAC were generated in simulation, in which all customers' batteries are operated in a purely self-interested manner and a counterfactual unconstrained undersea line current can be estimated;

3. The required inputs for the Shapley value calculation were generated using the actual and counterfactual data collected from steps 1 and 2 above, and Algorithm 2 is executed to calculate the normalised Shapley value;

4. The normalised Shapley value were used to disburse the TasNetworks budgeted amount, which was given by $1/kWh for all battery energy dispatched during the peak under the actual NAC operation.

5. The payments amounts were communicated to Reposit Power.

Due to the difficulties encountered in deploying Algorithm 1 (sample-based Shapley value approximation), deployment of Algorithm 2 alongside the NAC trial in live trials was not possible until February 2019. Nonetheless, the next two sub-sections focus on the payments made by Algorithm 2 over two peak events:

(i) during the Queen's Birthday Holiday 2018 period and

(ii) during a live trial in February 2019.

The first event is analysed in detail to show the effects on rewards of using the Shapley value, while the second demonstrates the use of the Shapley value in a live run of the NAC.

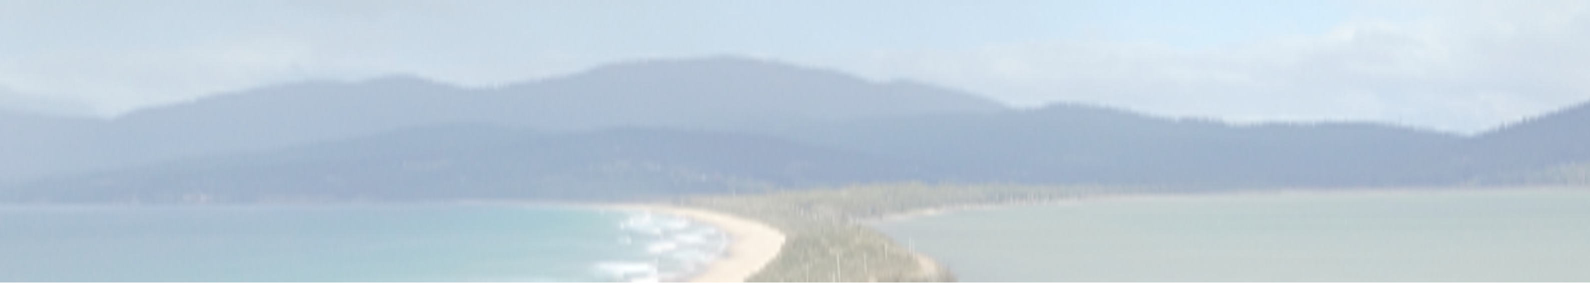

3.3.1 Example 1: Sunday 10 June 2018 morning peak

For this peak event, there were 31 batteries installed. Numbers of customers connected to each phase are given below:

Red-white	White-blue	Blue-red
5	19	7

This event was a true peak, and payments were made using the heuristic reward rule and *Energy Use* payment type described above (i.e. not the SV-based rewards); that is, the actual payments made by Reposit to the customers were calculate by the heuristic payment rule.

For comparison, the SV-based rewards presented here were calculated using Algorithm 2, where the sum of actual heuristic payments was used as the budget disbursed by the SV reward structure. The NAC outputs used in the SV calculations were either the actual NAC traces from the trial period, or a counterfactual run using a simulation of diesel dispatch with a human-in-the-loop.

The set of plots below illustrate the SV of the NAC network support activities on the morning of Sunday 10 June 2018. First, Figure 2 presents the line currents for each phase, which were the main driver of the SV-based payments.

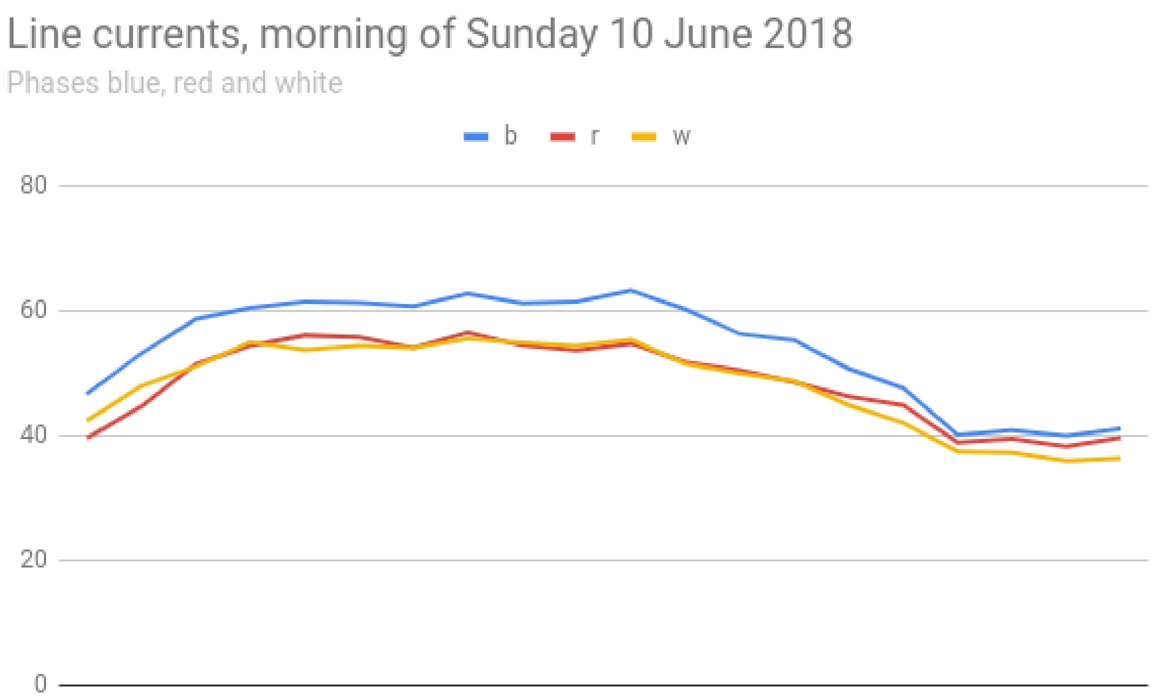

Figure 2. Line currents by phase on the undersea cable, morning of Sunday 10 June 2018, illustrating the thermal constraint exceedance on the blue phase, and phase unbalance.

Next, the Shapley value payments for the Sunday 10 June 2018 morning peak are plotted in Figure 3, coloured by phase-to-phase connection (blue is blue-red, red is red-white, and yellow is white-blue), and including the lower bound:

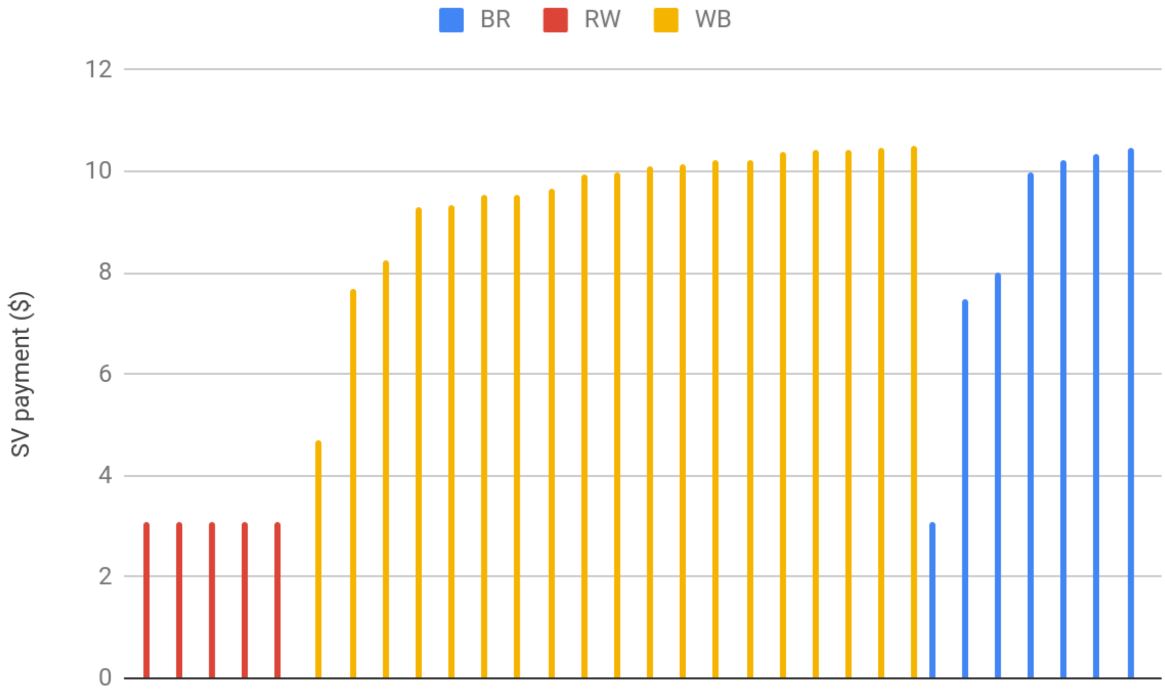

Figure 3. Shapley value payments by phase for the morning of Sunday 10 June 2018.

These plots show is that using the SV discriminates between the phases, by allocating greater payments to the batteries reducing the load on the line responsible for setting the dispatch level of the diesel generator (which is the blue phase). It is worth noting that such large differences in the Shapley value payments between phases is probably not likely to be a common occurrence in other networks. Here we have a combination of per-phase line current limits, large network imbalance, and low number of batteries relative to the unbalance. Modify any of these and the difference between phases would be quite small (see the second detailed example that follows). However, we conjecture that the biggest driver of Shapley value differences between customers grouped by phase connections will be whether there is a surplus or deficiency of available support between phases.

As shown in Figure 4, these SV payments are consistent with the heuristic, but do show some relative under-payment to the red-white phase connected customers.

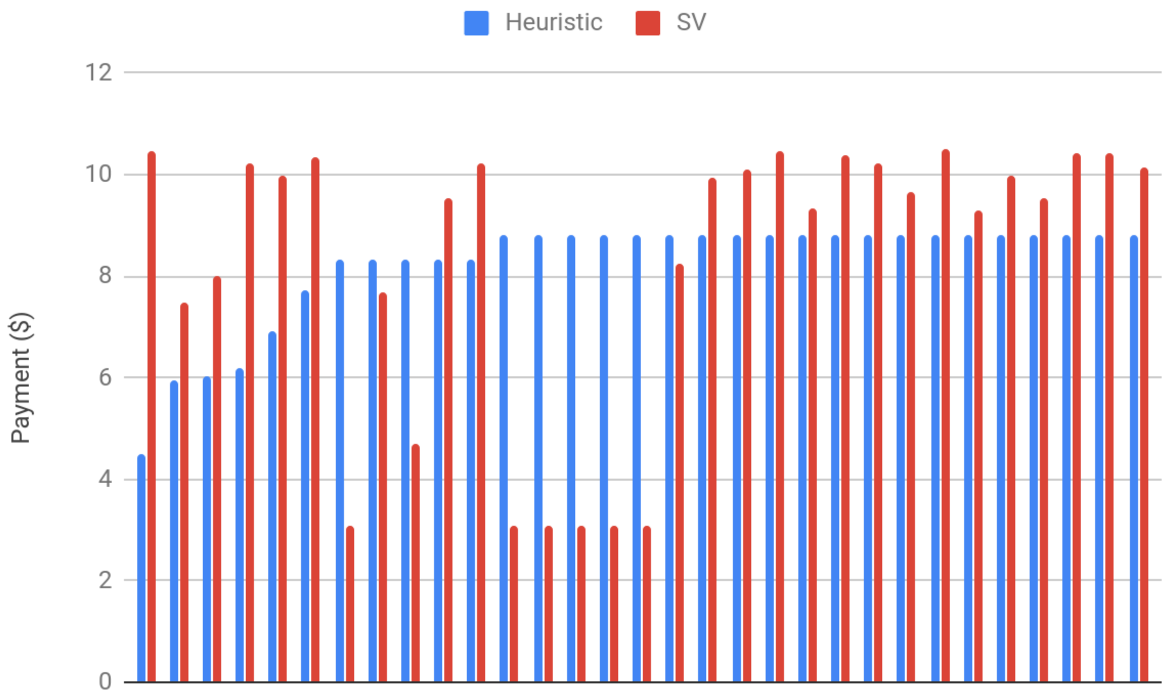

Figure 4. Heuristic (actual) payment vs Shapley value-based payment for the morning of Sunday 10 June 2018.

In part, this is an illustration of the shortcomings of the approximation when faced with large unbalance between the line currents. Specifically, it is only for a small number of counterfactual coalitions that either of the white or red phase line currents set the generator dispatch level. This small number of cases are missed when the customers' support is lumped by phase connection in the SV approximation. (The full extent of this effect is a key point of examination when the offline sample-based SV calculation is completed.)

The phase connections go a long way towards explaining the difference in SV payments between customers. However, other factors have considerable impact on the payments, and also can explain the divergence of the SV rewards and the heuristic payment ($1/kWh of energy dispatched from the battery during the peak period). The next two plots illustrate the effects of two of these factors, namely, voltage ratio and a customer's non-NAC counterfactual battery use.

Effect of voltage ratio: Figure 5 shows the relationship between SV reward and voltage ratio - the ratio of average diesel generator voltage to load connection point voltage during the peak event. Customer phase connection is indicated by colour; note that RW connections can be ignored because the SV is 0. For the remaining customers, this plot shows a slight upward trend. This is explained by noting that a lower voltage connection has a greater reduction of joule losses on the network than a higher voltage connection, and consequently less diesel use, for the same amount of energy discharged from the two batteries.

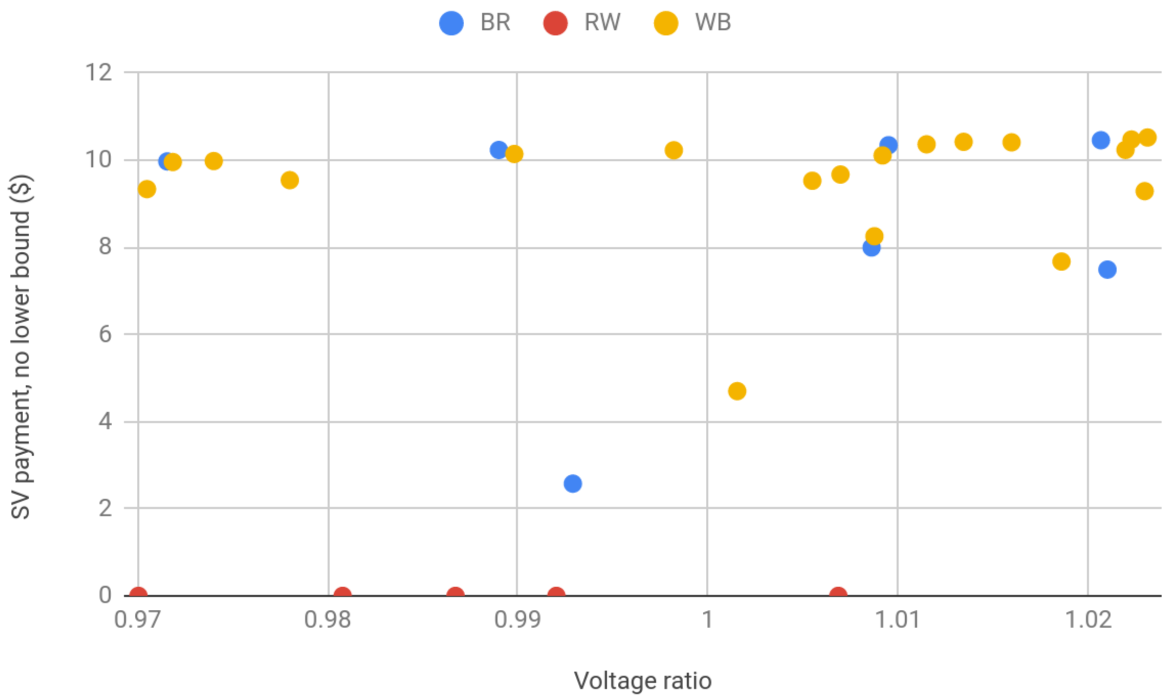

Figure 5. Voltage ratio vs Shapley value payment (Sunday morning, 10 June 2018).

Statistical analysis of this relationship is equivocal, mainly due to the small number of data points available. However, several customers on BR and WB connections deviate considerably from this trend. The next plot examines a second factor that explains part of this.

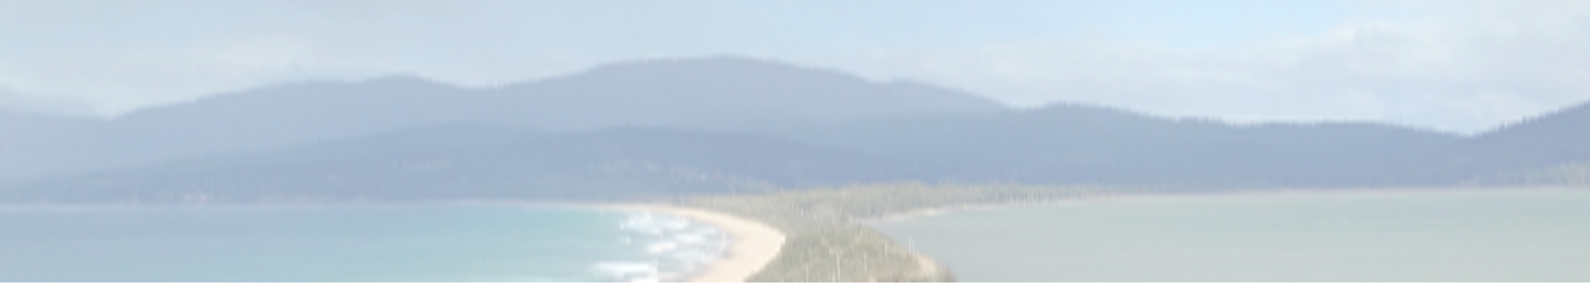

Effect of customers' non-NAC counterfactual battery use: Figure 4 illustrates the relationship between a customer's battery use in the non-NAC counterfactual and the Shapley value payment. Specifically, the customers' non-NAC counterfactual battery use is given by the depth of discharge during a simulation of the morning peak event without the NAC operating. The aim here is to measure how much the customer would have used their battery for their own purposes if the the NAC was not running. Again, customer phase connection is indicated by colour, and again RW connections can be ignored.

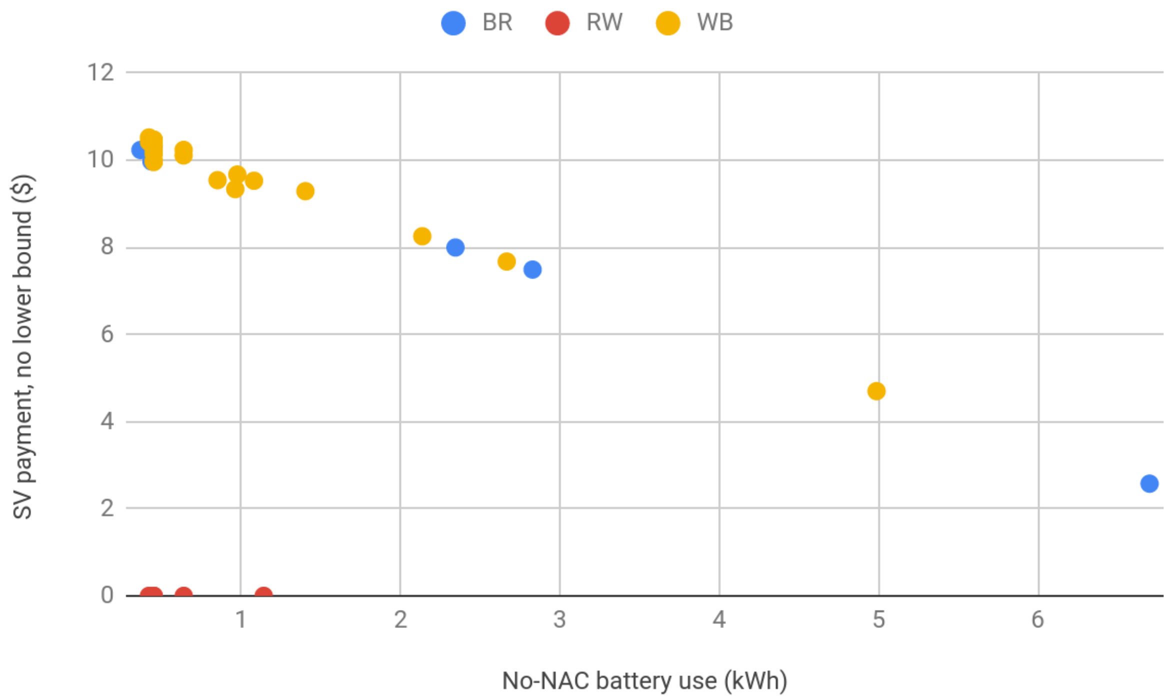

Figure 4. Battery use in the non-NAC counterfactual vs Shapley value payment (Sunday morning, 10 June 2018).

Figure 4 shows a strong negative relationship between non-NAC battery use and SV. In addition, the more extreme values on this plot, at ~$2.50 and ~$4.75, and the cluster around $8.00, correspond to the off-trend data points in Figure 3. This confirms that the SV reward structure rewards customers for only the additional benefit to the network that they bring for contributing to the NAC's operation, and not for discharging the battery within the peak period to service their own loads.

Although statistical analysis revealed a significant linear relationship, its interpretation is questionable due to the fact that this is a single peak event and additional significant variation in input data is expected between different peaks, the high likelihood of interacting terms in the model (e.g. phase connection and Shapley value), and the general problem of a low number of data points available from a single peak event. In addition, these data are generated using a heuristic that itself has considerable limitations. As such, the numerical statistical results are not presented here, because they could be both misleading and easily misinterpreted.

3.3.2 Example 2: Saturday 9 February 2019 morning (live trial)

Beginning early February, a number of synthetic peaks were induced on the network in order to test the computational routines that integrate Algorithm 2 (Shapley value approximation) with the NAC operation. This trial demonstrated an automated process of computing the Shapley value payments, and preparing them to be sent to Reposit Power.

Implementation: From a technology deployment point of view, these trials were a success, with all computational and data-handling procedures running as expected. specifically, these processes shadowed the continued use of the heuristic payment described at the start of Section 3.1, so were not used to compute actual payments to customers. Nonetheless, we now examine them as a test case for the Shapley value-based reward structures.

SV validation: The plot in Figure 5 below shows the *counterfactual* line currents (i.e. without NAC operation) assuming a reduced 40A line thermal limit. From this, it can be seen that the blue phase is carrying relatively more current and requires greater current injection to be reduced below the 40A limit. However, notice that all three phases require load reduction to suppress the diesel generator, and that the degree of unbalance between the phases is less than in the Queen’s Birthday case.

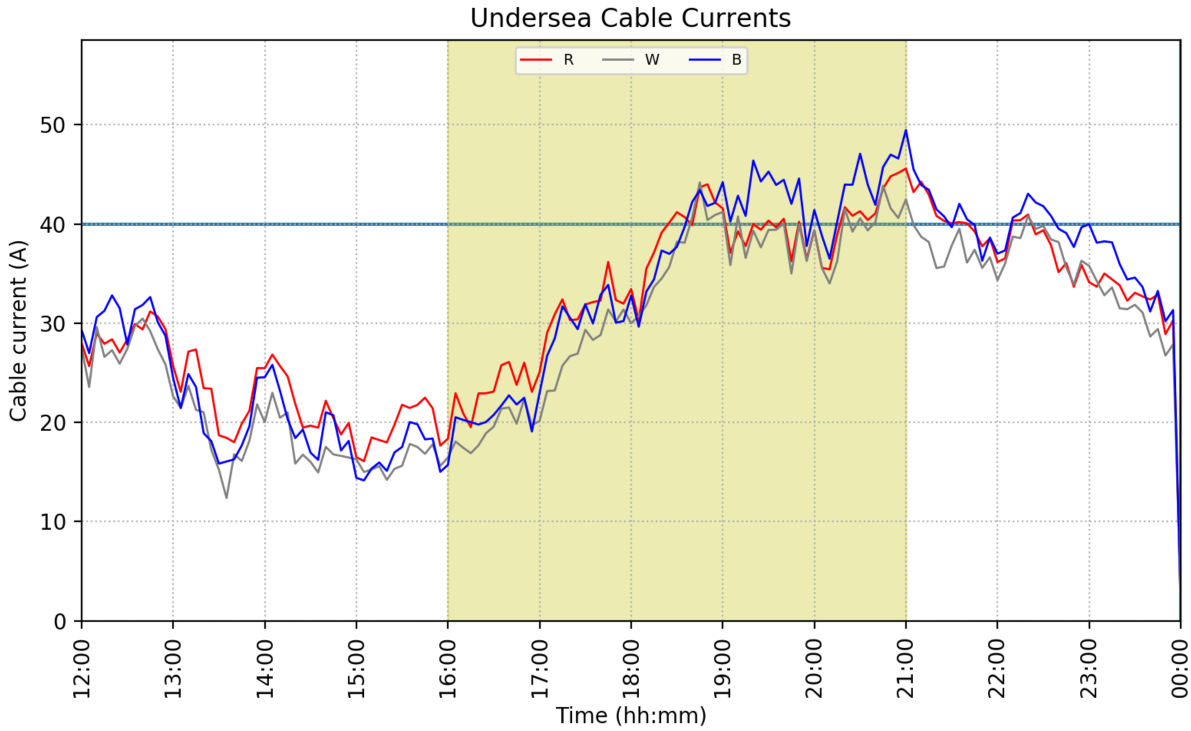

Figure 5. Counterfactual line currents for morning of Saturday 9 February 2019.

Next, in Figure 6 we compare the payments to each customer, arranged by phase (red are red-white connections, grey are white-blue, blue are blue-red). These Shapley value results (crosses, sv_support.json) were produced online, and here they are compared to the heuristic payment (circles, support.json) to the same customer (which was the payment actually made).

At an aggregate level, the payments to each phase under the heuristic and the Shapley value are in line with each other, in that the average values by phase are very similar. This is in keeping with the findings for the Sunday 10 June 2018 morning peak, in the section above. However, in this case, the payments for all customers are greater than zero. This is a reflection of the relatively small phase unbalance seen in this case study, compared to the earlier case.

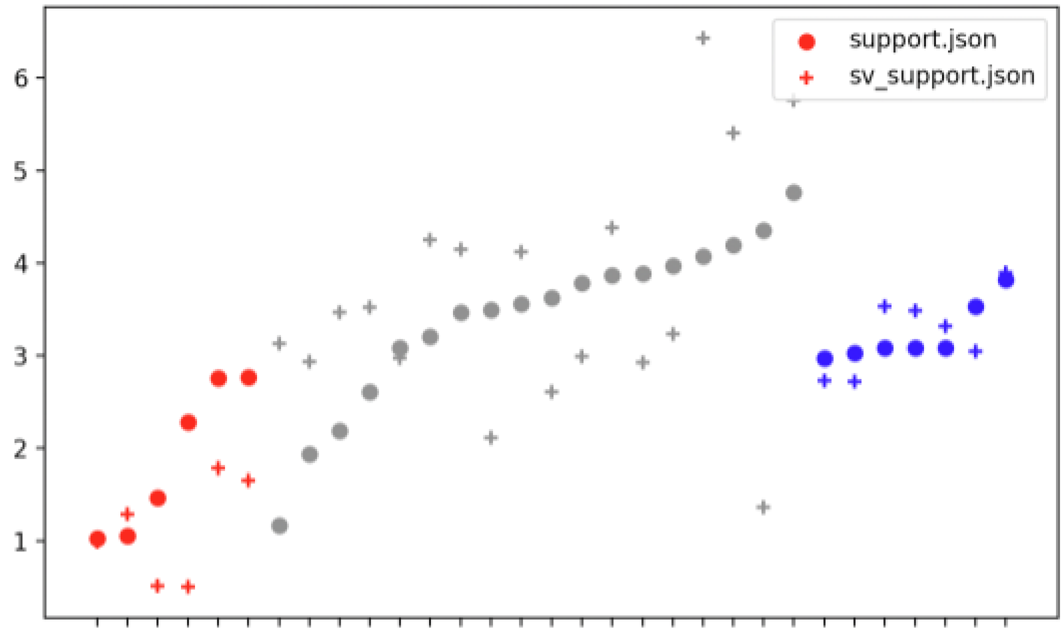

Figure 6. Heuristic and SV payments (Saturday 9 February 2019).

Deeper analysis of the effects of voltage ratio and a customer's non-NAC counterfactual battery use on the Shapley value rewards plotted here find evidence for the same relationships identified earlier, and again, they explain the bulk of the difference between heuristic payments and the SV payments.

3.3.3 Summary of case studies

The SV payments computed using Algorithm 2 reflect the desired properties of accounting for: 1) battery energy capacity and power limits, 2) line losses reductions, via the use of voltage ratios, and 3) customer's battery use in the non-NAC counterfactual. The results above show how well-designed reward structures can capture these effects, and show how support service pricing can be directly linked to the value that DER provide to regulated monopolies.

However, we should note that the required approximations in Algorithm 2 undermine the use of these reward structure for calculating customer payments in more complicated problems of sharing multiple DER value streams. In particular, the NAC was tested for simultaneously managing network voltages as well as thermal limits. In this instance, the approximations in Algorithm 2 were found to be poorly aligned with payment expectations, due to the difficulty of including both network management considerations (i.e. with non-linear network effects).

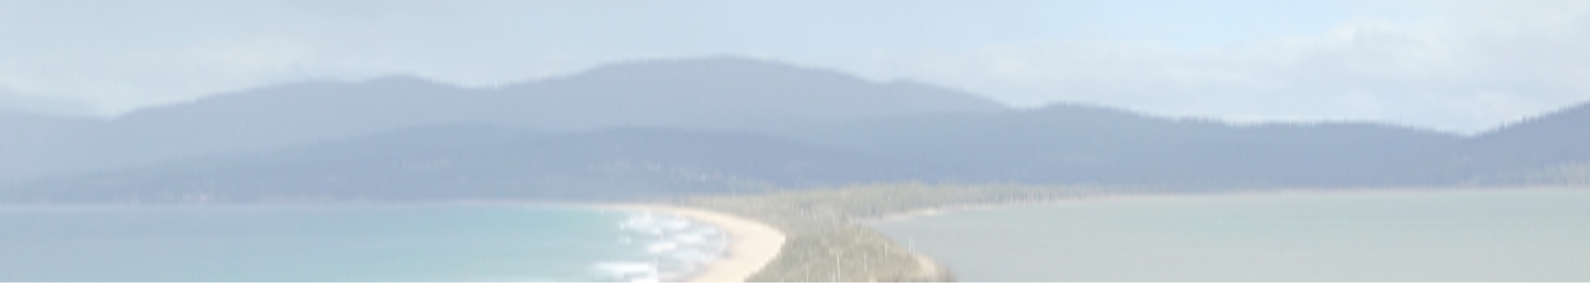

4 Customer financial analysis of network support payments

The aggregated rebates paid to customers for participation in the Bruny Island trial over the trial period March 2018- February 2019 are given in Table 2, below. Note that only 28 customers have participated over the full period, so this table includes some more recent additions to the Bruny Island fleet (those connected for <6 months are highlighted in cyan):

Customer ID	Total network support payments	Customer ID	Total network support payments
BT101	129.56	BT120	128.28
BT102	132.19	BT121	150.21
BT104	148.9	BT122	85.56
BT105	100.77	BT123	124.71
BT106	4.13	BT124	99.1
BT108	80.16	BT125	74.77
BT109	84.19	BT126	139.27
BT110	97.2	BT128	94.97
BT111	134.31	BT129	106.58
BT112	102.3	BT131	122.42
BT113	124.12	BT132	85.98
BT115	127.38	BT135	120.65
BT116	134.84	BT136	118.23
BT117	145.32	BT137	40.58
BT118	130.93	BT140	119.48
BT119	140.52	BT141	7.74

Table 2. Network support payments March 2018- February 2019

Note that these values reflect the heuristic payments under both the Energy Reserve and Energy Use payment types, at a flat price of $1/kWh dispatched during the peak periods. As such, network effects are not explicitly incorporated into the calculations. These payments, if repeated into the future, would provide cash-flows with net present values (NPVs) to customers given in the first column of Table 2 on the following page. The NPV values are the additional present value of cash-flows associated with running the NAC on Bruny Island under its current configuration (ie. with diesel generation), and are to be added on to the private benefits the customers earn for improving their rates of PV self-consumption.

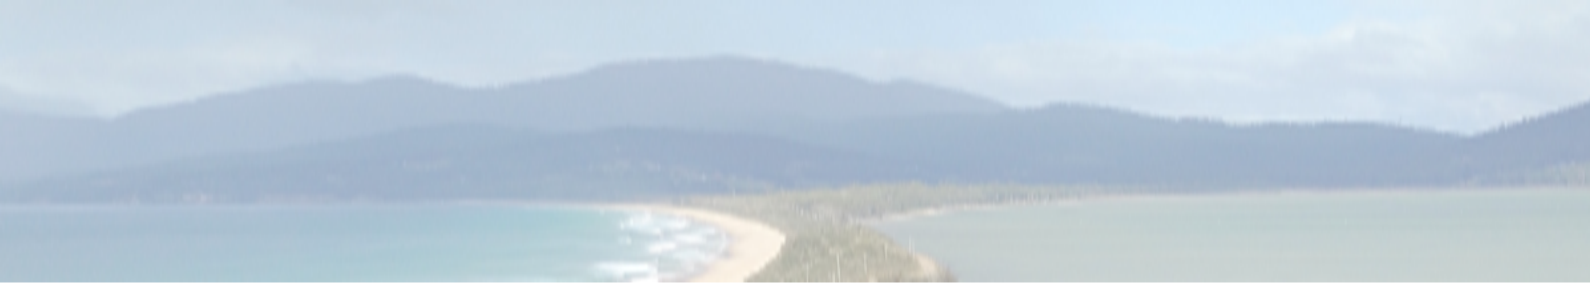

	Payment rate ($/kWh dispatched)			
Customer	**1**	**0.5**	**0.35**	**0.2**
BT101	$1,205	$603	$422	$241
BT102	$884	$442	$309	$177
BT104	$1,385	$693	$485	$277
BT105	$938	$469	$328	$188
BT106	$1,218	$609	$426	$244
BT108	$746	$373	$261	$149
BT109	$378	$189	$132	$76
BT110	$904	$452	$317	$181
BT111	$1,250	$625	$437	$250
BT112	$1,160	$580	$406	$232
BT113	$1,155	$577	$404	$231
BT115	$1,185	$593	$415	$237
BT116	$1,255	$627	$439	$251
BT117	$1,352	$676	$473	$270
BT118	$952	$476	$333	$190
BT119	$800	$400	$280	$160
BT120	$1,194	$597	$418	$239
BT121	$783	$392	$274	$157
BT122	$1,138	$569	$398	$228
BT123	$992	$496	$347	$198
BT124	$922	$461	$323	$184
BT125	$696	$348	$243	$139
BT126	$1,296	$648	$454	$259
BT128	$1,307	$654	$458	$261
BT129	$1,398	$699	$489	$280
BT131	$1,139	$570	$399	$228
BT132	$796	$398	$279	$159
BT135	$1,230	$615	$430	$246
BT136	$1,100	$550	$385	$220
BT137	$853	$426	$298	$171
BT140	$1,123	$561	$393	$225
BT141	$1,112	$556	$389	$222

Table 2. NPV for 15 year battery trial duration, discount factor 8.5%. Top row indicates payment reduction factors, left column of values is NPVs of actual payments made over 15 years.

The chance of external factors driving change in the associated network is quite high, reflected in the choice of discount rate at 8.5%, which is well above current WACC or finance rates.

The mean and median NPVs across all customers is approximately $1000, which may be a salient figure and provide a significant additional value stream to other customers considering investing in a residential battery. Note, however, the following caveats to the values reported here:

- The payments were made at $1.00/kWh discharged, subject to a very conservative lower bound on the payment level. This likely inflates the rebate paid to the customer over the economic value of the network support action, as measured by diesel fuel cost reductions. If fuel costs, as a measure of short run marginal costs, had been used to rebate customers, then this would be reflected in a price closer to $0.50/kWh. (Note: See the CONSORT social science report [2] for analysis of participants perspectives on the Network support payments. These were sometimes felt to be a “bonus” when compared to bill savings, so it is likely that network payments have a bit of room to move, when compared to retail tariffs, in the minds of households.)

- The payments were not intended to be generally applicable aggregate payments, but to reflect the situation on Bruny Island, where battery dispatch offsets diesel generator use. Different network settings may have lower economic values associated with the battery discharge activities, such as the value of capital deferral (cost-of-capital savings) or smaller capital works.

For this reason, several other NPV values are reported above, for network support payments made at $0.50, $0.35 and $0.20 per kWh dispatched. These figures are intended for discussion only, and may be reflective of other network investment or augmentation opportunities (not specified). Of course, it may be the that even greater value per kWh is possible for next-best network investment and operation options with greater costs.

A more detailed assessment of customer financial outcomes is presented in the Participants’ Solar and Battery System Financial Performance final report [3].

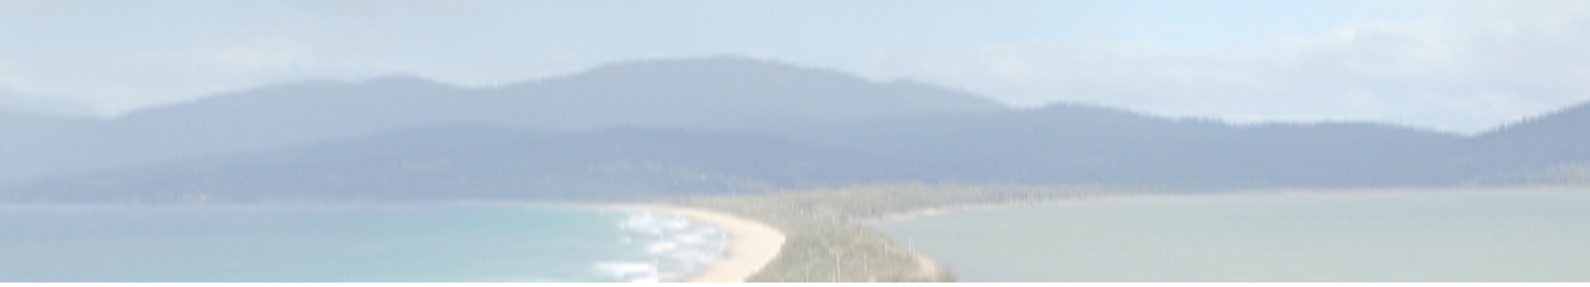

5 Conclusions

The CONSORT project addressed the Shapley value-based reward structures from an early stage of basic research to a deployed method of paying participants for use of their DER. Our results demonstrate how these methods implement value-reflective network support pricing for behind-the-meter distributed energy resources. The pricing methods developed can overlay any DER control scheme, from fully-fledged model-based optimisation methods like the NAC, down to much less sophisticated methods such as static control policies based on local voltage measurements, as all that is required is a simulation engine for the control scheme. We will continue to develop the required models and methods, and to prosecute the arguments, for value-based reward structures for support provided by customer-owned DER.

On the other hand, the difficulties encountered in bringing these method to live trials indicate the limits of its use for generating “spot” prices. However, it can be argued that network support is not well suited to spot market arrangements, and that longer-term posted prices, or even contracts for capacity, are more appropriate. Given this, the reward structures developed in the project can be used for generating such posted prices or contracts.

5.1 Recommendation: Methods for pricing network support agreements

In particular, the statistical analysis provided alongside the case study results present a path to using value-reflective reward structures for generating network and power system support services. Specifically, as an alternative to online “spot price” calculation, offline Shapley value computation, analysis and regression-based implementation maybe a viable way to implement the Shapley value-based methods for pricing network support.

The process could work as follows. Give a given a set of past network peak events, the Shapley value for each of them can be calculated offline (i.e. in simulation). The results would then be matched with appropriate input data, such as average in-peak load, battery use, or average voltage ratios. Regression of these inputs against the Shapley values would reveal the major contributing factors to the network support value, in simple to understand terms.

Such a regression-based model could, ultimately, take the characteristics of a home as inputs, and return the resulting prices specified in simple, but tailored, $/kWh units. The development of this type of tool would help both DNSPs and customers understand the mechanics of implementing value-reflective pricing of network support services.

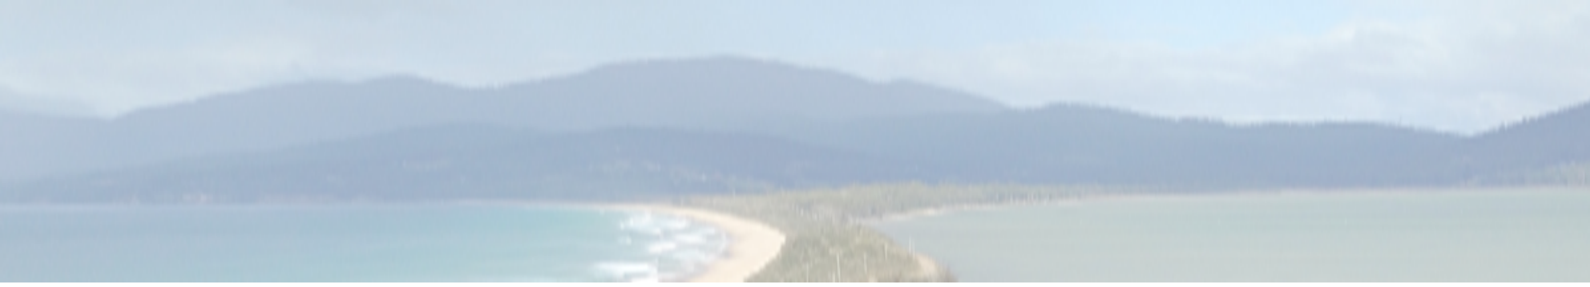

References

[1] Paul Scott, Dan Gordon, and Sylvie Thiébaux. Network-Aware Coordination, Final Report, CONSORT Bruny Island Battery Trial Project. April 2019.

[2] Phillipa Watson, Heather Lovell, Hedda Ransan-Cooper, Veryan Hann, and Andrew Harwood. Social Science, Final Report, CONSORT Bruny Island Battery Trial Project. April 2019.

[3] Evan Franklin and Archie Chapman. Participants' Solar and Battery System Financial Performance, Final Report, CONSORT Bruny Island Battery Trial Project. April 2019.