

Network-Aware Coordination of Residential Distributed Energy Resources

Paul Scott, Dan Gordon, Evan Franklin, Laura Jones, and Sylvie Thiébaux

Abstract—Rooftop solar and batteries, along with other distributed energy resources (DER), add a new demand-side flexibility, which, when harnessed, will enable distribution operators to more efficiently manage their constrained networks. This paper presents *network-aware coordination* (NAC), an approach for coordinating DER within unbalanced distribution network constraints, which utilises the alternating direction method of multipliers (ADMM) to solve a distributed receding-horizon OPF. As far as we are aware, this paper is the first to report on the practical implementation and performance of an ADMM-based technique solving a significant network problem in live operations. We present real-world trial results of NAC coordinating 31 residential batteries on a constrained feeder within Tasmania’s 11 kV distribution network. The batteries are coordinated to manage the network’s constraints during periods of high feeder demand, decreasing the need for expensive conventional network management, in this case a diesel generator. We achieve a 34% reduction in diesel over 7 peaks with 31 batteries capable of meeting 10% of peak demand. Supplementary simulations indicate the potential for a 74% diesel reduction if battery numbers were increased to 100. We find that compared to uncoordinated battery response, the NAC achieves 13% lower total costs over the trial period.

Index Terms—Distributed OPF, ADMM, Demand Response, DER, Prosumers, Batteries, Distribution Network, Live Trial

I. INTRODUCTION

Distributed energy resources (DER) and smart home energy management systems (EMSs) are transforming passive consumers of electricity into flexible *prosumers*, who can shape their energy use over time. This flexibility presents distribution network service providers (DNSPs) with opportunities to improve wider network performance, and offers a new cost-effective alternative to network augmentation and overly restrictive connection rules. However, many of these opportunities can only be realised if the actions of prosumers are coordinated. In fact, without coordination prosumers may cause more harm than good. The opportunities include:

- 1) successfully integrating large amounts of rooftop solar and electric vehicles (EVs) into our networks — without coordination, distribution networks may face increased risk of being overloaded or having voltage limits violated;
- 2) the management of existing ongoing network issues, shifting demand away from peak periods; and
- 3) demand-side participation in wholesale electricity markets in a way that respects the underlying limitations of the distribution network.

Paul, Dan and Sylvie are in the College of Engineering and Computer Science at the Australian National University, Canberra, Australia. Evan is in the School of Engineering at the University of Tasmania and Laura is in Network Innovation at TasNetworks, both in Hobart, Australia

As a first step, the electricity sector is seeing a rise in virtual power plants (VPPs)¹ and aggregators, that provide DNSPs with the opportunity to remotely control DER, in particular batteries. This is a key service, but the question remains as to *how* a DNSP should coordinate the DER on its network to meet the goals of keeping the network within its operating limits at low cost and high reliability.

The core goal of DER coordination can be effectively formulated as a multi-period optimal power flow (OPF) problem. Over the last 20 years, distributed optimisation, in particular the alternating direction method of multipliers (ADMM) [1], has been used to solve large-scale OPF problems [2], with attention more recently turning to using ADMM for DER coordination [3], [4]. Distributed optimisation has the potential to scale to large numbers of consumers, decouple DER owners and VPPs from the network operator, reduce the flow of private information, and increase resiliency.

Even with these benefits, the underlying OPF problem for coordinating DER is still very challenging to solve, and distributed optimisation faces significant challenges in real-world operations. These challenges include: uncertainty in network models and customer loads; modelling real distribution feeder components; communication overhead; and solving in a timely, online manner. To date it has not been demonstrated whether an ADMM-based approach can be designed and implemented to overcome these real-world challenges, and furthermore, produce results that show a tangible benefit over competing approaches available to network operators.

We address this gap in the literature, by designing, implementing and field trialling an ADMM-based approach that we call *network-aware coordination* (NAC). The NAC incorporates a unique combination of techniques that make it work in a real setting: an online receding-horizon, unbalanced 3-phase network modelling and a simple decoupling of the participants from the network. We use ADMM to solve a series of distributed multi-period OPF problems online in a receding horizon manner, guided by a series of network-level and residential forecasts. The OPF contains a detailed unbalanced 3-phase network model, which can correctly account for the per-phase contribution of DER. Finally, our decomposition at only the connection between the houses and the network strikes a practical balance between enabling parallelisation of the ADMM subproblems, and limiting the number of ADMM iterations required to converge — which is particularly important when real-world communications delays come into effect.

¹Companies such as Reposit Power, Evergen, Tesla and Sonnen.

Taken individually, these techniques have previously been proposed for use with ADMM, but as far as we are aware, they have not been integrated and experimented with in a real setting. Our second contribution is to demonstrate and evaluate our approach’s performance in real-world trials solving a real network problem. This provides important insights into the feasibility and performance of these ADMM-based approaches, that will help to guide their further development and refinement. As far as we are aware, trial results have not been presented before for online distributed OPF approaches for coordinating residential DER.

The trials were conducted on an 11 kV feeder on Bruny Island, Tasmania, which has a binding line capacity during peak periods. The trials were run for 65 days, managing 16 peak periods with a combination of 31 residential battery systems and a diesel generator. We quantify the diesel savings for all 16 peaks, and provide more detailed counter-factual simulations for 7 of the peaks to quantify the performance relative to alternatives available to the DNSP.

The next section of this paper compares our approach to existing literature. We then present the Bruny Island feeder as a motivating network problem. The NAC approach is presented in four stages: the receding horizon optimisation, the multi-period unbalanced OPF solved for each horizon, the ADMM decomposition of the OPF, and finally an overview of the implementation. The remaining sections of the paper explain the on-network trials and accompanying simulations, and present the results and conclusions.

II. RELATED WORK

The NAC approach builds on prior work, including our own [3], for solving OPF problems in a distributed manner using ADMM. In general terms, the literature has varied in the level of network model detail, whether it is a transmission or distribution system, which network-connected devices are coordinated, how they decompose the OPF problem, what additional technologies the ADMM approach integrates with, and the sophistication of the experiments. Motivations consist of coordinating the actions between independent microgrids [5], [6], decomposing the network to gain greater parallelisation of subproblems [2], [7]–[9], providing participants with greater agency and privacy [3], [4], and establishing complete decentralisation where messages only are passed between locally connected network components [10], [11].

As stated in the introduction, while the individual techniques that go into our NAC approach have been proposed before in an ADMM context, it is their unique combination which has made the NAC work in practical setting. In the following, we explain where these techniques have been used individually in related work, and the differences to our approach.

A receding horizon (also known as model predictive control) was proposed for use with decentralised OPF in [3], [4]. Here we produce a working implementation, along with online forecasts that enable us to assess its ability to manage uncertainty.

Phase unbalance can play a significant role in many distribution network problems, with individual customers and their DER having single phase connections to the network [8].

ADMM has been used to solve OPF problems in an unbalanced setting in [8], [12], [13]. These focus on more conventional single-period OPFs, while we solve a multi-period OPF to schedule DER, which greatly increases the problem size. In addition to this new use, the abstract multi-phase network model we have developed presents a flexible framework for constructing and solving 3-phase network problems. It enables the network components such as open-delta transformers to be straight-forwardly modelled, something that can require significant effort in approaches based around an admittance matrix [14]. The abstraction also makes it possible to apply different power flow models, and the explicit representation of connections makes it easy represent different decompositions.

Compared to most related work, in particular [3], [4], [10], [11], we take a more reserved approach to decomposing the problem in these trials by doing so only at the customer-network boundaries. While more aggressive forms of network decomposition can lead to greater parallelisation of the subproblems, they may also lead to an increase in the number of iterations required before the problem converges. These extra iterations can quickly become a significant bottleneck when real-world communication delays and data limits are taken into account. That said, our approach and model have been designed to enable this extra level of decomposition if it proves beneficial for particular distribution network feeders, the trade-offs of which we expect to explore in the future [15].

As far as we are aware, this is the first work to demonstrate online distributed OPF coordinating residential DER in real-world trials. The literature review in [16] provides a good overview of existing demand-side flexibility field studies, here we highlight some of the more relevant studies.

The Olympic Peninsula Project [17] managed a variety of DER including residential electric water heaters, municipal water pumps and distributed generators. They had access to approximately 75 kW for the price-responsive homes. The problem they solve is a much simpler unit dispatch problem, while we model power flows within the feeder itself, solving a full OPF that accounts for network losses and phase contributions. Their trial had to use a *virtual* feeder to demonstrate managing feeder capacity, because in reality the resources were spread across several feeders. Our trial is over a real feeder solving a real network constraint problem.

The Pacific Northwest Smart Grid Demonstration Project [18] trialled a transactive energy system where energy and prices were exchanged between 27 nodes at the transmission level. Our trial instead focuses on network problems down at the distribution level, where we solve an OPF problem using distributed optimisation.

A different transactive energy approach, named Power-Matcher, was trialled in the PowerMatching City project [19]. This approach uses agent bidding and market clearing techniques, but again does not account for distribution power flows and hence does not solve an OPF.

III. THE BRUNY ISLAND FEEDER

We present the Bruny Island feeder as a motivating example of how DER coordination can provide an improved outcome,

compared to uncoordinated DER operation or conventional constraint mitigation. Bruny Island is a small island off the southern coast of Tasmania, Australia, with over 1000 dwellings. Electricity is supplied to the island from mainland Tasmania via two undersea 11 kV cables. The southern cable, shown in Figure 1, supplies over 800 customers via a long rural feeder. It 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, the DNSP responsible for the network, currently manages this problem by locating and dispatching a portable 550 kVA diesel generator unit on the island [20].

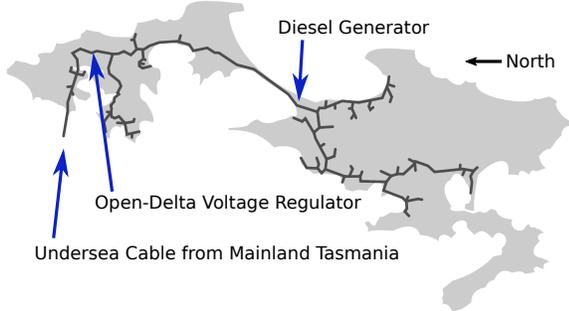


Figure 1: Bruny Island southern MV feeder.

The southern feeder has around 337 buses (once reduced), a peak load of around 1.3 MVA, and experiences significant phase unbalance which varies with time, with up to around 15% difference in current between phases.

This work was conducted as part of the CONSORT² project: a research and industry collaboration that is trialling the use of around 31 customer-owned battery storage systems to help TasNetworks manage network problems on Bruny Island. The NAC platform was developed and used to coordinate Reposit Power³ EMS equipped battery systems in live trials throughout 2018 and 2019. There are two other key components to the project not explored in this paper: the social science of this new customer, network and technology relationship [21], and the game theory and financial analysis of new ways to reward customers for their DER support [22].

While the trials are restricted to this one network feeder, the approach itself was designed to be able to expand beyond Bruny Island, to other feeders with existing or emerging issues associated with DER deployment, to coordinate additional DER beyond batteries, and to enable new interactions in wholesale markets. As such, many of the findings relating to performance, modelling and practical implementation will apply beyond Bruny Island.

IV. RECEDING HORIZON OPTIMISATION

Every 5 minutes the NAC solves (using ADMM) a multi-period OPF problem. Each multi-period OPF makes up one horizon in this receding horizon approach. The horizons cover up to 24 hours⁴, each made up of 48 time steps. With a

reoptimisation every 5 minutes, only the actions in the first time step of each horizon get acted on — the remaining steps are there to ensure that the decisions are not short-sighted.

Figure 2 illustrates seven consecutive horizons, with time steps aligned to the half-hour. This half-hour alignment is used to simplify integration with forecasting and to align to TOU tariffs. The first step can vary in duration between 5 and 30 minutes, depending on where the horizon lands on the clock, but it is always only ever acted on for 5 minutes.

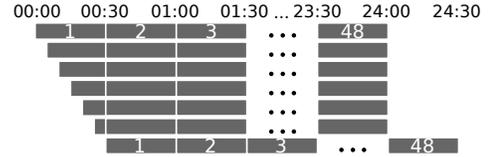


Figure 2: Seven consecutive horizons (top to bottom), starting from midnight, each made up of 48 time steps.

The optimisation for a horizon is initiated 4 minutes 15 seconds prior to the horizon start, which is the time available to find a solution. This leaves enough time (45 sec) for the forecasters to collect their latest inputs and update.

V. OPTIMAL POWER FLOW

Here, we describe the multi-period unbalanced OPF problem that is solved within each horizon. We begin by formulating an abstract model of a multi-phase network, defining how each of the components connect. As explained in Section II, this model is designed to be flexible so that new components can be easily integrated, and to make the decomposition simpler in Section VI. The OPF problem for this network over a single horizon is then presented, followed by definitions for the network elements used for the Bruny Island network, including lines, generators and houses with battery storage.

A. Abstract Multi-Phase Network Model

A network is a tuple (N, E, T, c, a) , where $N = \{n_1, \dots, n_{|N|}\}$ is a set of single-phase nodes, $E = \{e_1, \dots, e_{|E|}\}$ is a set of elements (lines, generators, etc.), and $T = \{\tau_1, \dots, \tau_{|T|}\}$ is a set of terminals (multi-phase points of connection for an element). The connection function $c : T \rightarrow N^{<N}$ connects each terminal to a sequence of nodes, where the notation $N^{<N}$ represents the set of finite sequences with entries from the set N . Finally, the association function $a : E \rightarrow T^{<N}$ associates an element with a sequence of terminals (e.g., one for load, two for line).

Figure 3 shows an example network with two single-phase loads, one supplied by a delta-wye transformer with neutral, and the other from a 2-wire spur. An approximate single line diagram is shown along with our explicit network representation that models the phases as separate nodes and sequences the connections between terminals (solid squares) and nodes (solid circles).

We use the shorthand $c_\tau := c(\tau)$ for terminal τ , and further $c_{\tau,k}$ gives k -th connected node for the terminal (the k -th entry in the sequence $c(\tau)$).

²See our website for more project details: <http://brunybatterytial.org>

³<https://repositpower.com>

⁴It is likely that, for this type of online use, shorter horizons could be used without much performance impact, e.g., [23] found 16 hours was sufficient.

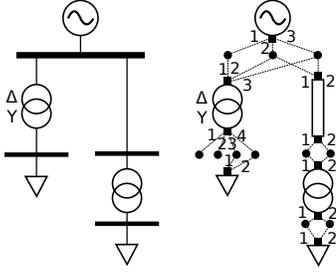


Figure 3: Left: single line diagram. Right: our explicit representation with nodes (solid circles), elements, terminals (solid squares) and ordered connections for each terminal (numbered dashed lines).

B. The Optimal Power Flow Problem

At the network level there are two main types of variables: potential variables and flow variables. Each node $n \in N$ has a potential variable $v_n \in V$. The terminals T are partitioned into three types: line current terminals T_I , line power terminals T_L , and line-to-line power terminals T_{LL} . Each terminal τ has one or more flow variables, for $k, k' \in \{1, \dots, |c(\tau)|\}$:

- $i_{\tau,k} \in I$ if $\tau \in T_I$
- $s_{L,\tau,k} \in S$ if $\tau \in T_L$
- $s_{LL,\tau,k,k'} \in S$ if $\tau \in T_{LL}$

In order to apply a flow conservation constraint, i.e. Kirchhoff's junction law (KJL), at each node we need to convert the different types of terminal variables into compatible quantities. We choose to convert them all to equivalent line currents by introducing new auxiliary current variables that represent the line current, line power and line-to-line power contributions. For all nodes $n, m \in N$:

$$i_n = \sum_{(\tau,k) \mid \tau \in T_I, c_{\tau,k}=n} i_{\tau,k} \quad (1)$$

$$i_{L,n}^* v_n = \sum_{(\tau,k) \mid \tau \in T_L, c_{\tau,k}=n} s_{L,\tau,k} \quad (2)$$

$$i_{LL,n,m}^* (v_n - v_m) = \sum_{\substack{(\tau,k,k') \mid \tau \in T_{LL}, \\ c_{\tau,k}=n, c_{\tau,k'}=m}} s_{LL,\tau,k,k'} \quad (3)$$

These are then summed to enforce KJL:

$$i_n + i_{L,n} + \sum_{m \in N} (i_{LL,n,m} - i_{LL,m,n}) = 0 \quad (4)$$

Each element $e \in E$ has objective and constraint functions:

$$f_e : X_e \rightarrow \mathbb{R}, \quad g_e : X_e \rightarrow \mathbb{R}^{d_{g,e}} \quad (5)$$

were $X_e = V^{d_{v,e}} \times I^{d_{i,e}} \times S^{d_{s,e}} \times Y_e$ captures the variables of the element, consisting of the potential variables of nodes it connects to, its terminal flow variables and some auxiliary variables internal to the element Y_e . The $d_{*,e} \in \mathbb{N}$ parameters capture the required variable and constraint dimensions for the element. The function g_e represents any hard constraints for element e , which are satisfied when $g_e(x_e) \leq 0$ (using an element-wise inequality). This is a general form for representing the constraints of elements — in the sections that follow we write the constraints in more a readable form.

The optimisation problem for minimising the sum of element objective functions (maximising social welfare) is:

$$\min_{v_n \forall n \in N, x_e \forall e \in E} \sum_{e \in E} f_e(x_e) \quad (6)$$

$$\text{s.t. } g_e(x_e) \leq 0 \quad \forall e \in E \quad (7)$$

(1) – (4)

C. Element Models

For the trials conducted in this paper, the potential and flow variables are vectors of complex numbers that represent the network voltage and current phasors, and complex power in rectangular form over a horizon of $|H|$ time steps: $V, I, S \in \mathbb{C}^{|H|}$. For example, $v_{n,t}$ is the voltage of node n at time step $t \in H = \{1, \dots, |H|\}$. The vector $\delta \in \mathbb{R}_{>0}^{|H|}$ provides the durations of the time steps. Here we only need elements with one or two terminals, so we simplify the notation by dropping the terminal index for an element's first terminal, and using a prime to distinguish its second: e.g., $i_{k,t} := i_{\tau,k,t}$ and $i'_{k',t} := i_{\tau',k',t}$, where $\tau, \tau' = a(e)$. Similarly, when the element needs to access the voltages its terminals connect to we use the notation: $v_{k,t} := v_{c_{\tau,k},t}$ and $v'_{k',t} := v_{c_{\tau',k'},t}$.

1) *Bus*: Bus elements can be used to conveniently group together nodes that are co-located, or to impose appropriate line-neutral or line-line voltage constraints. The flow variables for a bus are zero. For a 3-phase bus with line-line voltage lower and upper bounds $\underline{v}, \bar{v} \in \mathbb{R}_{\geq 0}$, the constraints are:

$$\underline{v}^2 \leq |v_{k_1,t} - v_{k_2,t}|^2 \leq \bar{v}^2 \quad \forall (k_1, k_2) \in \binom{\{1, 2, 3\}}{2} \quad (8)$$

$$i_{k,t} = 0 \quad \forall k \in \{1, 2, 3\} \quad (9)$$

Note that the voltage lower bound constraints are non-convex.

2) *Line*: An α -phase line has two current terminals which each connect α nodes at either end. We use a $2\alpha \times 2\alpha$ nodal admittance matrix $Y \in \mathbb{C}^{2\alpha \times 2\alpha}$ to model each line. By putting the terminal voltages and currents into vectors at each time step, $v_t := [v_{1,t}, \dots, v_{\alpha,t}, v'_{1,t}, \dots, v'_{\alpha,t}]$ and $i_t := [i_{1,t}, \dots, i_{\alpha,t}, i'_{1,t}, \dots, i'_{\alpha,t}]$, the line constraint is simply:

$$i_t = Y v_t \quad (10)$$

A thermal limit $\bar{i} \in \mathbb{R}_{>0}$ is imposed on each line current: $|i_{k,t}|^2 \leq \bar{i}^2 \quad \forall k \in \{1, \dots, \alpha\}$.

3) *Feeder*: A feeder has a single current terminal with 3-phases that represents the zone substation at the root of the feeder. We assume that there is a voltage regulator that does a perfect job keeping a steady voltage $\hat{v} \in \mathbb{C}^3$ over time (perfect voltage source). The constraints simply fix the connected nodes to these voltages, and leave the terminal currents free to take on any value: $v_{k,t} = \hat{v}_k \quad \forall k \in \{1, 2, 3\}$.

4) *Open-Delta Voltage Regulator*: An open-delta voltage regulator consists of the open-delta transformer shown in Figure 4 and a tap controller that maintains the secondary winding voltages. It has two current terminals.

We relax the regulator model by treating the taps as continuously adjustable between their limits. This is a reasonable approximation for our feeder, since each tap position translates

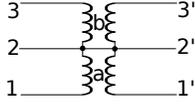


Figure 4: Open-delta transformer with two independent taps a and b. The primary side (first terminal) is on the left.

to a $\sim 0.6\%$ voltage change. The ratios for the taps a and b are modelled by the variables $r_{a,t}, r_{b,t} \in [\underline{r}, \bar{r}]$. We assume we have direct control over the tap ratios rather than relying on the regulator's inbuilt controller, and we assume an ideal transformer model. Even with the stated simplifications, the equations remain non-convex:

$$v_{2,t} = v'_{2,t} \quad (11)$$

$$(v_{1,t} - v_{2,t}) = r_{a,t}(v'_{1,t} - v'_{2,t}) \quad (12)$$

$$(v_{3,t} - v_{2,t}) = r_{b,t}(v'_{3,t} - v'_{2,t}) \quad (13)$$

$$i_{1,t} + i_{2,t} + i_{3,t} + i'_{4,t} + i'_{5,t} + i'_{6,t} = 0 \quad (14)$$

$$i'_{1,t} = -r_{a,t}i_{1,t}, \quad i'_{3,t} = -r_{b,t}i_{3,t} \quad (15)$$

5) *Fixed Loads*: These loads have a single terminal of type line current, line power or line-to-line power depending on the desired model. These terminal variables are fixed to a constant value for each time step to match the desired load. In this work we only require line-to-line fixed power loads.

6) *Generator*: We make use of a simple delta connected generator model that injects an equal amount of power into each phase. It has a single line-to-line power terminal. The generator can have a running cost and a minimum operating point (although we relax this in our experiments). For a 3-phase generator:

$$p_t + iq_t = s_{1,2,t} = s_{2,3,t} = s_{3,1,t} \quad (16)$$

$$\underline{p}z_t \leq p_t \leq \bar{p}z_t, \quad \underline{q}z_t \leq q_t \leq \bar{q}z_t, \quad z_t \in \{0, 1\} \quad (17)$$

$$f(\dots) = \sum_{t \in H} \left(\frac{1}{2} \Psi_t (3p_t)^2 + 3\psi_t p_t \right) \delta_t \quad (18)$$

where we have modelled the generators costs with a quadratic where $\Psi_t \in \mathbb{R}_{\geq 0}$ and $\psi_t \in \mathbb{R}$ are prices (e.g., with units \$/kW/kWh and \$/kWh respectively).

7) *House*: A house has a single line or line-to-line power terminal. In this work we connect them line-to-line to the MV network. As we do not have reliable LV network data, this line-to-line connection accounts for their influence on the MV network through a delta-wye distribution transformer.

A house is modelled as having a fixed background power consumption combined with a number of DER. In this work each house has a battery and solar PV system. The specifics of the household battery optimisation are handled by Reposit Power's EMS, which has some performance and behavioural tweaks for the benefit of customers. We do not present the full details, but instead a representative model that closely captures the key battery optimisation behaviour.

Here we assume a single phase house (line-to-line between two nodes) with a standalone battery inverter (not a hybrid inverter). The house has an uncontrollable background power consumption $s_{back,t} \in \mathbb{C}$ and battery power $s_{batt,t} \in \mathbb{C}$.

The single line-to-line terminal power / house connection point power is: $s_{1,2,t} = s_{back,t} + s_{batt,t}$. We impose a apparent power limit on the battery and inverter system: $|s_{batt,t}|^2 \leq \bar{s}_{batt}^2$. As is common, we relax the battery real power into separate charge and discharge components:

$$s_{batt,t} = (p_{c,t} - p_{d,t}) + iq_{batt,t} \quad (19)$$

A simple linear relation for the battery state of charge $SOC_t \in [0, \overline{SOC}]$ where we have charge and discharge efficiencies $\eta_c, \eta_d \in (0, 1]$ is then:

$$SOC_t = SOC_{t-1} + (\eta_c p_{c,t} - p_{d,t} / \eta_d) \delta_t \quad (20)$$

The objective function of the house consists of the retailer tariffs applied to their connection point power. Typically, different tariffs are offered for energy consumption and energy export, and only active power is typically metered. For $p_t + iq_t = s_{1,2,t}$, and letting γ_t^+, γ_t^- be the consumption and feed-in tariffs, the house's objective function is:

$$f(\dots) = \sum_{t \in H} \begin{cases} \gamma_t^+ p_t \delta_t & \text{if } p_t \geq 0 \\ \gamma_t^- p_t \delta_t & \text{if } p_t < 0 \end{cases} \quad (21)$$

This is relaxed under the condition that $\gamma_t^+ \geq \gamma_t^-$ by introducing auxiliary variable β_t :

$$f(\dots) = \sum_{t \in H} \beta_t \delta_t, \quad \beta_t \geq \gamma_t^+ p_t, \quad \beta_t \geq \gamma_t^- p_t \quad (22)$$

D. Model Summary

To summarise, the combined objective function consists of the diesel generator operating costs (18) and the sum of the participating household retail tariff costs (22). The model has non-convex constraints in the KJL power relations (2), (3), which occur for all nodes that connect to fixed load and house network elements. The open-delta voltage regulator has non-convex constraints where the tap ratios multiply the voltages and currents in (12), (13), (15). The lower voltage limit in (8) is a final source of non-convexity. Section VII discusses the implications of these non-convexities.

VI. ADMM DECOMPOSITION

This section provides a high-level overview of the ADMM algorithm which is used to decompose and solve the OPF problem for each horizon (for further details of the algorithm and its application to OPF see [1] and [3]).

We decompose the problem at the interface between the participating households and the network. The flow variables associated with a houses single power terminal are the only common variables between the house and the rest of the network. The general approach in ADMM is to duplicate such common variables, keeping one for the network $s_t = p_t + iq_t$, and one for the house $\hat{s}_t = \hat{p}_t + i\hat{q}_t$, and to link them together with a constraint $s_t = \hat{s}_t$. An augmented Lagrangian relaxation is taken of this constraint:

$$\mathcal{L} = \sum_{t \in T} \left(\frac{1}{2} \rho (p_t - \hat{p}_t)^2 + \frac{1}{2} \rho (q_t - \hat{q}_t)^2 \right) + \lambda_{p,t} (p_t - \hat{p}_t) + \lambda_{q,t} (q_t - \hat{q}_t) \quad (23)$$

where the λ terms are dual variables / Lagrangian multipliers for the constraint, and ρ is a penalty parameter. The dual variables are locational marginal prices (LMPs) for power with units \$/kW and \$/kVar for real and reactive power, which incentivise the house EMS to act in a way to support the network when it becomes constrained. This augmented Lagrangian gets added to the objective function for each house, therefore the EMSs can account for these side-by-side with existing time-of-use (TOU) and feed-in tariff (FIT) arrangements they have with their retailer. As described in [24] and [25], these distribution-level LMPs can form the basis of a market for managing network constraints (we use them within a receding horizon rather than day-ahead market).

The ADMM algorithm has three stages per iteration k :

- 1) Obtain s_t^k by minimising over network variables, using λ_t^{k-1} and holding \hat{s}_t constant at \hat{s}_t^{k-1} .
- 2) Obtain \hat{s}_t^k by minimising over house variables, using λ_t^{k-1} and holding s_t constant at s_t^k .
- 3) Update the dual variables: e.g., $\lambda_{p,t}^k = \lambda_{p,t}^{k-1} + \rho(p_t^k - \hat{p}_t^k)$

The implication is that we can decompose our large OPF into much smaller subproblems that are solved iteratively to come to an overall solution. In the first stage of the algorithm, the network subproblem only needs to know the terminal power of the houses from the previous iteration — it does not need to know any of the underlying household constraints, auxiliary variables or parameters. Furthermore, as all of the time coupling constraints (those associated with the battery state of charge) only appear in house subproblems, the network subproblem can be split up into many smaller subproblems and solved independently for each time step, i.e. by solving $|H|$ unbalanced 3-phase OPF problems in parallel.

For the second stage, each house can solve its own small subproblem independent of the network constraints and all other houses. In addition to enabling greater parallelisation, this decomposition presents a clear separation of responsibility and ownership between the VPPs and DNSPs, and enables multiple, potentially competing, VPPs to operate over the same network at the same time.

VII. NAC IMPLEMENTATION

Our NAC implementation has three main modules: a Dealer, Workers (network solvers), and Remotes (EMSs or VPPs). The Dealer keeps track of time and initiates the optimisation for a horizon by sending jobs to Workers, which solve the distribution network part of the problem, and to Remotes, which solve the EMS part of the problem. The Dealer updates prices and checks for convergence of the ADMM algorithm.

The flow of data between these modules is shown in Figure 5. A Redis⁵ store is used for all network, forecast and participant connection data. Redis is also used as an inter-process communication layer, with the Dealer pushing jobs onto a queue, and Workers popping these jobs. This enables the system to easily scale to multiple CPU cores (simply spawn more workers) and across multiple machines. HTTP is used for asynchronously communicating with remotes, using JSON payloads to initiate a Remote computation.

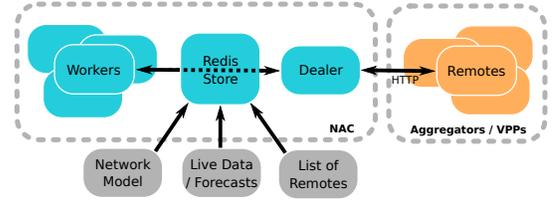


Figure 5: Structure of the NAC implementation and interaction.

The Workers use the Madopt⁶ interface to Ipopt [26] to model and solve the NLP network problems, in these experiments using the HSL [27] ma27 library as a backend. As discussed in Section V-D, the underlying network problem for the Workers is non-convex. We use Ipopt directly on this non-convex problem, which will return the first locally optimal solution it finds. In the single phase case, previous work has shown experimentally that often these local solutions are within 1% [3], [7] of the global optimal. We leave investigating the global optimality gap for the 3-phase unbalanced case as a future task — this will require developing tight convex relaxations for the unbalanced case with an open-delta transformer.

A. Forecasters

Online forecasts are critical for the receding horizon approach. We split the forecast uncertainty into two sources: that from network customers that do and do not participate in NAC. The participant Reposit Power EMSs have their own personalised load and solar PV forecasts. The non-participant forecast is made at an aggregate feeder level, and then proportionally allocated as loads at the distribution transformers.

Existing feeder-level load forecasting options, such as those currently built into a DNSP's SCADA or DMS, were not suitable for our online receding-horizon forecasting needs. They were either designed for much larger network sections or rigid day-ahead usage. We decided to train our own forecasters for the trial, with a simple linear regression model as a first attempt. An improved transformer neural network (TNN) forecaster was developed, which we describe in [28]. The forecaster takes as input recent recloser load readings and air temperatures from a weather station on Bruny Island. The output is the island load at half-hourly resolution over a forward horizon of 24 hours.

B. Infrastructure

The Workers, Dealer and forecaster were running on a server with two 6 core Intel Xeon CPUs (L5640 @ 2.27 GHz, circa 2010). This server communicates with Reposit Power who run a (lightly utilised) cloud 2-node cluster (each with 2 cores, 4 GB memory) to solve the Remote subproblems. Once the negotiations for a horizon have converged, the decisions are sent out to participant EMSs where they are then acted on. To save on cloud costs, the Remote subproblems could instead be solved directly by the EMSs on participant premises. However, due to unreliable and slow internet connectivity on the island, this approach was deemed too much risk for the project.

⁵<https://redis.io>

⁶<https://github.com/stanle/madopt>

VIII. TRIALS AND SIMULATIONS

A. Trial Details

To date, the NAC has been trialled in live operations for a total of 65 days. These cover 5 trial periods where peaks are known to have historically occurred [20]. The trial periods and the number of peaks they experienced that required either diesel or battery support are:

- March 29 to April 3 2018 (6 peaks)
- April 13 to May 3 2018 (0 peaks)
- June 8 to June 12 2018 (5 peaks)
- July 13 to July 23 2018 (5 peaks)
- December 19 to January 9 2019 (0 peaks)

Only those particular days where peaks occurred, and hence required DER coordination, are of interest here.

Various improvements were made to the NAC implementation between the March and June trials. Further improvements were made on July 16 when the forecaster was switched from the linear regression to the TNN approach. For these reasons, the more detailed analysis focuses on the results in the later trial periods: 3.5 days covering 5 peaks on June 8–11 and 2 days covering 2 peaks on July 20–21. The NAC diesel savings are analysed for the other 9 peaks in Section IX-A.

During the trials we elected to reward participants for their network support at a generous fixed rate of \$1/kWh on the amount their battery was discharged during peak events. The customer payments could be quite uncertain if directly exposed to NAC LMPs, as explored in Section IX-B, and hence it was deemed too high risk to do so for the trials to date without first gaining a better understanding of their volatility.

Nevertheless, while this might have social and economic implications, it does not impact the results presented in this paper. The EMSs were still responding correctly to the LMPs in the background, the participants just get rewarded with more generous payments after the event.

B. Simulation and Model Details

There were relatively few peaks available for live trials, and the network conditions during these peaks were uncontrollable and varied. To provide crucial counter-factual comparisons to the status quo and alternative battery coordination techniques, we therefore supplemented the live results with a number of simulations:

- 1) Simulation replaying NAC (NAC Sim)
- 2) NAC with perfect island load forecast (NAC Perf)
- 3) Manually scheduled dispatch strategy (Manual)
- 4) Independent local battery optimisation (Indep)
- 5) Participants without batteries (No Batts)
- 6) NAC with 100 batteries (NAC 100)

Simulations were performed using the open source Smart-GridToolbox simulation library⁷. In this section we discuss some of the points of difference between the simulation and the real-world.

The connection phasing of some participants was unknown, so had to be randomly allocated. We estimate around 75%

of them to be correct in the simulation. The simulation does not replay any of the communications problems that were experienced during the trials. During the live trials, several of the participants had communications outages at various times.

In the NAC, the zone substation was modelled as a constant voltage source (feeder element), and part of the network on mainland Tasmania between the zone substation and undersea cable was removed. Simulations over the July 20–21 period show that this simplification results in a maximum relative error in the current through the undersea cable of 2.6%, or less than 1 A; most of the time, the error is substantially smaller. This, along with further network reductions reduced the number of feeder buses from 517 to 337.

These model inaccuracies and the large load ramp rates around the peak periods present a challenge to the receding-horizon OPF approach which only makes a new decision every 5 minutes (the cable current can change by up to 6 A in 5 minutes). Our approach has been to set a more conservative cable line limit in the NAC model, which means that it might at times provide more support than is strictly required. As discussed below, the diesel is typically started when the cable current reaches 64 A, while for most of our trials we have used a 57 A limit in the NAC model (recent improvements have allowed us to increase it to 60 A).

A human operator manually controls the diesel generator by reacting to real-time SCADA alarms that trigger on the cable current. We approximately model this in the simulations, using a “human-in-the-loop” (HITL) controller that mimics the response a human operator could achieve if they were paying attention to the alarms and reactively adjusting the diesel output to prevent a cable line limit violation.

Briefly, the simulated HITL controller responds to a smoothed version I_s of the measured cable line current I_c (maximum across phases), specified by the following equation at the j -th time step:

$$I_{s,j} = I_{c,j} \left(1 - e^{-\Delta t/\tau}\right) + I_{s,j-1} e^{-\Delta t/\tau} \quad (24)$$

where Δt is the time step and τ is a smoothing time constant. Sufficient power is then injected at the diesel generator so as to maintain this smoothed current at or below the cable limit, while enforcing a lower bound on the power that the diesel generator can supply. For the simulations in this paper, we used $\tau = 9$ minutes, a minimum diesel power of 100 kW and a cable limit of 64 A, which were tuned to match good historical operator responses.

IX. RESULTS

Figure 6 shows the aggregate battery and diesel response for 5 trial days (June 9–11 and July 20–21), alongside the participating house load and solar output. The shaded green regions indicate the presence of non-zero NAC LMPs, i.e. times where the NAC believes the undersea cable constraint is binding. For most peak periods both battery and diesel response is required, except the evening peak on July 21 where the batteries managed the cable constraint on their own. The NAC conservatively dispatches the batteries a small amount

⁷<https://github.com/NICTA/SmartGridToolbox>

on the morning of the two July days. The counter-factual simulations suggest that without this support the cable constraint would have been near its limit, but not quite violated. We therefore do not label these periods as peaks, but still count the cost of these conservative NAC actions.

The first simulation (NAC Sim) was run as a point of reference between the real-world trials and the simulations. It replays the measured conditions from the trial but with simulated network, EMS actions, and HITL controller. Its results are overlaid on Figure 6 for comparison to the trial.

Over both trial periods the mean absolute error (MAE) between the real and simulated aggregate battery power is 7.8 kW. The real human operated diesel control and our HITL controller have a MAE of 32.7 kW if we focus just on the 6 peak periods where the diesel ran. This is not surprising, as a human operator in the control room has many other tasks to take care of, so cannot always devote as much attention to managing the diesel response as our HITL controller.

This is most observable in the July 20 evening peak in Figure 6. The operator dispatched the diesel more than was necessary to meet the cable constraint during this peak, and left the diesel running for more than an hour after the peak had subsided. As such, the diesel usage was two times more than necessary when compared to the HITL controller. This illustrates that more automated control of the diesel generator on its own could achieve significant cost savings.

A. Costs

Figure 7 shows the costs that the NAC achieved in reality alongside the simulations. These are split up into the diesel component (on top) and aggregate participant component, which is mostly related to their retail tariffs. The overall objective is to minimise the sum of these two. Figure 8 provides June 10 as an example comparison between the aggregate battery and diesel response for the different approaches (the power direction follows load convention). This day has both morning (0600) and evening (1800) peaks.

The actual trial (NAC) has a 7% higher cost relative to the simulation (NAC Sim). This is a result of the diesel generator being left on for the July peak discussed in the previous section (an extra \$73), and once accounted for the simulation agrees with the live trial total costs to within 0.7%.

Understandably, the case where there are no batteries in the system (No Batts) has the largest cost. When batteries are added without coordination (Indep), the participants obtain a 17% benefit (by locally minimising their costs) and incidentally also reduce diesel usage. Figure 8 shows the batteries do not act for the morning peak in the Indep case (no precharging of the battery) but they do contribute a small amount to relieving the evening peak.

The Manual dispatch is a hand-tuned schedule for the batteries. The rule, aggregate dispatch of 88 kW from 08:00 to 10:50 and 53 kW from 16:30 to 21:15 on each day, was developed based on historical data on the peak timings, intensity and durations. It achieves a good reduction in diesel usage, but it over-utilises the batteries and hence significantly increases participant costs.

The NAC achieves a 13% reduction in costs compared to the uncoordinated Indep case. The participants incur a slightly higher cost so that the diesel can be significantly reduced. As such we would expect the network to use some of these savings to compensate the participants for their assistance. 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 an up to 3.5% improvement with a better forecast. The morning peak in Figure 8 demonstrates the impact of a forecast that underestimates the upcoming peak, which leads to not enough battery precharging. While the receding horizon approach enables the NAC to correct as it gets closer to the peak, it still only manages half the dispatch of the perfect forecast during the peak.

The 31 batteries in the trial, around 4% of all feeder customers, have the capacity to meet 10% of peak feeder demand, which managed to achieve a 34% reduction in diesel use compared to the case with no batteries. To see how this changes with the number of batteries, a NAC simulation with 100 batteries was conducted, which reduced the diesel consumption usage by 74%.

Using the same counter-factual analysis, the NAC was found to have achieved a 30% reduction in diesel for the 6 peaks over the March 29 trial period (when the approach was less refined) and 60% reduction in diesel over 3 peaks from July 13. The 3 peaks in this July period are relatively small compared to other peaks, which is why it achieves a higher percentage reduction.

B. Prices

Figure 9 shows histogram plots for the non-zero real-power NAC LMPs experienced by each participant in the June period. The width of the line indicates the relative frequency at which the given customer experienced the price on the y-axis. The customers are first sorted based on their network phasing (5 RW, 7 BR, 19 WB) and are then further sorted within these groups based on their distance from the zone substation. There is a clear pattern of higher prices on the most common phase (WB), which typically experiences the most load. There is also a pattern of higher prices offered at the end of the feeder, where a reduction in load can avoid more network losses.

As discussed in Section VIII-A, customers were in reality paid at a fixed rate of \$1/kWh for the amount their battery was discharged. If they had instead been paid according to the LMPs (applied to the house connection point power rather than just battery power), the mean payment to customers during the June trial would have been \$3.88. This corresponds to a total net payment of \$120 from the network to customers, which in aggregate is more than enough to compensate customers for their increased tariff-related costs incurred by providing network support. The maximum payment to a customer would have been \$14.6, and the minimum payment -\$8.1 (network charges customer due to large local demand). In fact, 9 of the 31 customers would have overall had to pay the network if LMPs were enforced, which could deter participation in a voluntary scheme, in the absence of additional incentives.

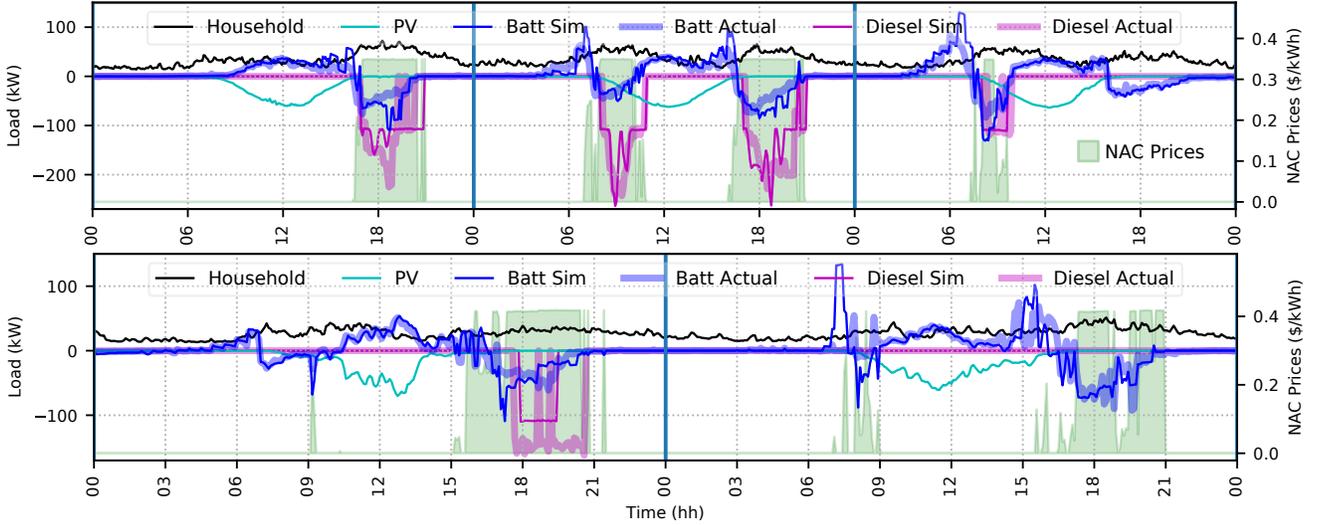


Figure 6: Battery and diesel response for five days (June 9–11 top, July 20–21 bottom) comparing the actual and simulated values.

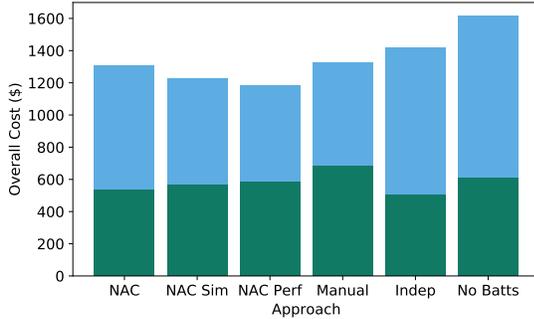


Figure 7: Diesel (upper segment) and participant (lower segment) costs for different approaches over the 5.5 days for detailed analysis.

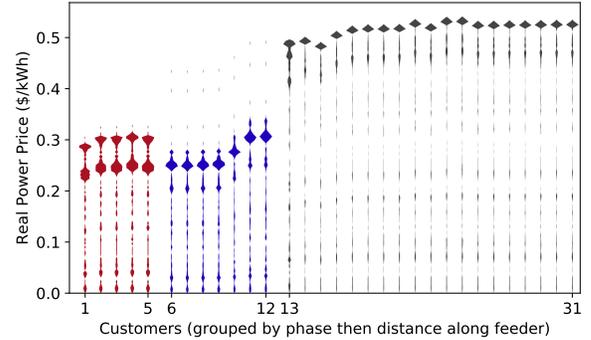


Figure 9: Histogram of real power LMPs for each customer.

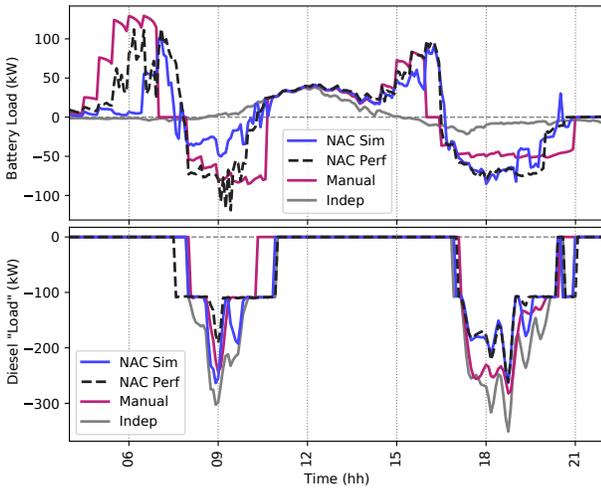


Figure 8: Aggregate battery (top) and diesel (bottom) power compared for June 10.

C. Computational Performance

ADMM was used with a penalty of $\rho = 0.1$ \$/kVA². Two sets of tolerances were selected, an *acceptable* tolerance of 2×10^{-3} for the primal and dual residuals, and a *desired* tolerance of 5×10^{-4} . The negotiation continues until it reaches the

desired tolerance. If time runs out, then the EMSs will only act if an iteration met the acceptable tolerance.

For the June trial the algorithm reached the desired tolerance in 960 out of 982 horizons, with the remaining 22 solved to acceptable values. The worst case real power disagreement (or residual) between the network and customers was 8 W across all customers, time steps and horizons that met the desired tolerance. Over the remaining 22 “acceptable” horizons the worst case was 80 W, and the mean was less than 1 W.

Over all horizons the mean number of iterations was 18.7 with a standard deviation of 15, and the maximum was 64. The average time per iteration was approximately 4 seconds, while 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.

The equivalent simulations (NAC Sim) produced somewhat better convergence for the June period, with a mean of 13.6 iterations per horizon. During the live trials, an EMS’s predictions could be updated mid-negotiation, slowing down convergence, while the simulation presents a more stable environment. In the 100 battery experiments the mean number of iterations is actually slightly better at 13.0. The mean

Worker solve time goes up slightly to 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.

After factoring in the utilisation of the infrastructure, the extra computational and communications costs required to run the NAC are estimated to be \$4.5 per day where support is required. This reduces the benefit of the NAC over the Manual dispatch to 5.6%.

X. CONCLUSION

We have presented NAC for coordinating residential DER on a distribution feeder. The online receding horizon approach and decomposition has enabled distributed multi-period OPF to be solved in a practical setting. We have demonstrated the ability of such an approach to reduce diesel consumption on a real-world trial, and outperform alternatives, advancing the technology closer to the point where it presents a credible tool for DNSPs to manage their networks in a future with high levels of renewables and other DER.

Future experiments will focus on scaling to thousands of participants, enabling the reactive power response capabilities, exploring the use of soft and hard voltage constraints, and establishing the value of LV network modelling.

Uncertainty remains a big challenge to the reliability of these techniques. This can be tackled this from two different angles: improved general-purpose load and solar forecasting techniques for use in an online setting; and exploring enhanced techniques such as distributed robust optimisation and coupling scheduling with fast, intelligent local control.

ACKNOWLEDGEMENTS

The Australian Government, through the Australian Renewable Energy Agency, is providing \$2.9m towards the \$8m CONSORT project under its Research and Development Programme. Thanks to all the CONSORT project partners at the Australian National University, TasNetworks, Reposit Power, the University of Sydney and the University of Tasmania.

REFERENCES

- [1] 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>
- [2] B. Kim and R. Baldick, "A comparison of distributed optimal power flow algorithms," *Power Systems, IEEE Transactions on*, vol. 15, no. 2, pp. 599–604, 2000.
- [3] P. Scott and S. Thiébaux, "Distributed Multi-Period optimal power flow for demand response in microgrids," in *ACM e-Energy*, Bangalore, India, jul 2015, "Best Paper Award".
- [4] M. Kraning, E. Chu, J. Lavaei, and S. Boyd, "Dynamic network energy management via proximal message passing," *Foundations and Trends in Optimization*, vol. 1, no. 2, 2014.
- [5] Y. Liu, H. B. Gooi, and H. Xin, "Distributed energy management for the multi-microgrid system based on ADMM," in *2017 IEEE Power Energy Society General Meeting*, July 2017, pp. 1–5.
- [6] W. J. Ma, J. Wang, V. Gupta, and C. Chen, "Distributed energy management for networked microgrids using online ADMM with regret," *IEEE Transactions on Smart Grid*, vol. 9, no. 2, pp. 847–856, March 2018.
- [7] T. Erseghe, "Distributed optimal power flow using ADMM," *IEEE Trans. on Power Systems*, vol. 29, no. 5, pp. 2370–2380, Sept 2014.
- [8] E. Dall'Anese, H. Zhu, and G. B. Giannakis, "Distributed optimal power flow for smart microgrids," *Smart Grid, IEEE Transactions on*, vol. 4, no. 3, pp. 1464–1475, 2013.
- [9] Q. Peng and S. H. Low, "Distributed Algorithm for Optimal Power Flow on a Radial Network," *ArXiv e-prints*, Apr. 2014.
- [10] P. Šulc, S. Backhaus, and M. Chertkov, "Optimal Distributed Control of Reactive Power via the Alternating Direction Method of Multipliers," *ArXiv e-prints*, 2014.
- [11] S. Magnússon, P. C. Weeraddana, and C. Fischione, "A Distributed Approach for the Optimal Power Flow Problem Based on ADMM and Sequential Convex Approximations," *ArXiv e-prints*, Jan. 2014.
- [12] B. A. Robbins and A. D. Domínguez-García, "Optimal reactive power dispatch for voltage regulation in unbalanced distribution systems," *IEEE Trans. on Power Systems*, vol. 31, no. 4, pp. 2903–2913, July 2016.
- [13] H. J. Liu, W. Shi, and H. Zhu, "Distributed voltage control in distribution networks: Online and robust implementations," *IEEE Transactions on Smart Grid*, pp. 1–1, 2017.
- [14] R. Yan, Y. Li, T. K. Saha, L. Wang, and M. I. Hossain, "Modeling and analysis of open-delta step voltage regulators for unbalanced distribution network with photovoltaic power generation," *IEEE Transactions on Smart Grid*, vol. 9, no. 3, pp. 2224–2234, May 2018.
- [15] J. Guo, G. Hug, and O. K. Tonguz, "Intelligent partitioning in distributed optimization of electric power systems," *IEEE Transactions on Smart Grid*, vol. 7, no. 3, pp. 1249–1258, May 2016.
- [16] P. Kohlhepp, H. Harb, H. Wolisz, S. Waczowicz, D. Müller, and V. Hagenmeyer, "Large-scale grid integration of residential thermal energy storages as demand-side flexibility resource: A review of international field studies," *Renewable and Sustainable Energy Reviews*, vol. 101, pp. 527 – 547, 2019. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1364032118306944>
- [17] D. J. Hammerstrom, R. Ambrosio, J. Brous, T. A. Carlon, D. P. Chassin, J. G. DeSteese, R. T. Guttromson, G. R. Horst, O. M. Järvegren, R. Kajfasz, S. Katipamula, L. Kiesling, N. T. Le, P. Michie, T. V. Oliver, R. G. Pratt, S. E. Thompson, and M. Yao, "Pacific northwest gridwise testbed demonstration projects: Part 1. olympic peninsula project," Pacific Northwest National Laboratory, Tech. Rep., 2007. [Online]. Available: <https://eio.pnnl.gov/research/gridwise.stm>
- [18] R. Melton, "Pacific northwest smart grid demonstration project technology performance report volume 1: Technology performance," 2015. [Online]. Available: <https://www.osti.gov/servlets/purl/1367568>
- [19] F. Bliet, A. van den Noort, B. Roossien, R. Kamphuis, J. de Wit, J. van der Velde, and M. Eijgelaar, "Powermatching city, a living lab smart grid demonstration," in *2010 IEEE PES Innovative Smart Grid Technologies Conference Europe (ISGT Europe)*, Oct 2010, pp. 1–8.
- [20] E. Franklin, D. Gordon, D. Jones, P. Scott, L. Blackhall, and S. Thiébaux, "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.
- [21] H. Lovell, V. Hann, and P. Watson, "Rural laboratories and experiment at the fringes: A case study of a smart grid on bruny island, australia," *Energy Research & Social Science*, vol. 36, pp. 146 – 155, 2018, spatial Adventures in Energy Studies:. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S2214629617303158>
- [22] A. C. Chapman, S. Mhanna, and G. Verbič, "Cooperative game theory for non-linear pricing of load-side distribution network support," in *The 3rd IJCAI Algorithmic Game Theory Workshop*, 2017.
- [23] P. Scott, S. Thiébaux, M. van den Briel, and P. Van Hentenryck, "Residential demand response under uncertainty," in *International Conference on Principles and Practice of Constraint Programming (CP)*, Uppsala Sweden, sep. 2013, pp. 645–660.
- [24] L. Bai, J. Wang, C. Wang, C. Chen, and F. Li, "Distribution locational marginal pricing (dlmp) for congestion management and voltage support," *IEEE Transactions on Power Systems*, vol. 33, no. 4, pp. 4061–4073, July 2018.
- [25] R. Li, Q. Wu, and S. S. Oren, "Distribution locational marginal pricing for optimal electric vehicle charging management," *IEEE Transactions on Power Systems*, vol. 29, no. 1, pp. 203–211, Jan 2014.
- [26] A. Wächter and L. T. Biegler, "On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming," *Mathematical Programming*, vol. 106, no. 1, pp. 25–57, 2006. [Online]. Available: <http://dx.doi.org/10.1007/s10107-004-0559-y>
- [27] HSL Archive, "A collection of fortran codes for large scale scientific computation," 2018. [Online]. Available: <http://www.hsl.rl.ac.uk>
- [28] 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.