Accounting for Forecast Uncertainty in the Optimized Operation of Energy Storage

Khalid Abdulla, Kent Steer, Andrew Wirth, Saman Halgamuge
Melbourne School of Engineering, University of Melbourne
Melbourne VIC 3010, Australia
Email: kabdulla@student.unimelb.edu.au

Julian de Hoog
IBM Research – Australia,
19/60 City Road,
Melbourne VIC 3006, Australia

Abstract—**This paper presents and empirically evaluates two approaches to accounting for forecast uncertainty when attempting to optimize the operation of a residential battery energy storage system. Data-driven methods are used for forecasting, and dynamic programming, within a receding horizon controller, is used for operational optimization. The first method applies a discount factor to costs incurred at later intervals in a deterministic dynamic programming control horizon, provided with point forecasts. In the second approach probabilistic (scenario) forecasts are generated using Lloyd-Max quantization of the distribution of forecast errors, to allow the use of a stochastic dynamic programming formulation. These methods are applied to maximizing the cost-savings delivered from a residentially owned and operated battery, using a case-study of residential consumers with roof-top PV systems in New South Wales, Australia. It is found that scenario forecasts can offer an 8% increase in annual cost-savings, on average, when using a univariate multiple linear regression forecast.**

I. Introduction

There is increasing interest in the role which distributed energy technologies, such as residential energy storage, embedded generation, and demand response, can play in the transformation of the way we generate, transmit, store and consume electrical energy. Distributed control of these technologies is attractive for two reasons; it avoids the need to install and maintain extensive communication infrastructure, and it maximizes the resilience offered by a more distributed electrical infrastructure.

For distributed control to be most effective, it is often necessary to forecast future values of, for example, the demand of a single customer, or small collection of customers, or the output of a single roof-top PV system. It is well known that using data-driven approaches to forecast demand and PV output are more effective at large scales, and that forecasts for small-scale aggregations suffer from large forecast errors [1], [2]. As a result it is helpful to consider this forecast uncertainty when attempting operational optimization of distributed energy assets.

This paper is concerned with operational optimization of a battery energy storage system which is owned by, and operated in the interests of, a residential customer. Under the tariff structure assumed, which is typical of residential customer tariffs in Australia, there are two ways in which a battery can be operated to reduce the electricity costs of a customer who must satisfy a local demand, and has a roof-top PV (Photo-voltaic) system. The first is solar-self-consumption, *i.e.* minimizing the export of excess PV generation to the grid which is rewarded a price much lower than the import cost. The second is time-of-use tariff optimization, *i.e.* minimizing grid imports during peak-price times. To maximize the cost-saving offered by a battery, it is necessary to forecast both the local demand, and the PV output, to allow good charging decisions to be made.

There is an extensive literature of methods for performing optimization under uncertainty, Sahinidis [3] provides a broad and relatively recent review of state-of-the-art methods, including operational optimization under forecast uncertainty, which is of interest in the present study. The operational optimization of an energy storage system is inherently a multi-stage problem, and lends itself to solution using Dynamic-Programming (DP), with the battery state-of-charge offering a natural state variable. As a result DP is a popular choice for energy storage operational optimization [4]–[8].

To be of practical interest, operational optimization must be formulated in a way that can be applied in an on-line setting. A natural way of doing so is to formulate the problem as a receding horizon controller. The present authors have previously formulated a receding horizon controller to optimally operate a residentially owned battery, with each horizon solved using dynamic programming [9].

Formulating the problem as a receding horizon controller immediately offers some resilience to forecast errors, as control decisions further into the control horizon (where forecasts are likely to be less certain), are subject to recourse. However, it is likely that performance can be improved if forecast uncertainty is explicitly considered when finding the horizon-optimal solution. In this paper we apply two methods in an attempt to improve the operational optimization of energy storage, subject to forecast uncertainty. The following contributions are made in terms of optimizing the operation of energy storage:

(i) A discount factor is applied within the DP solution

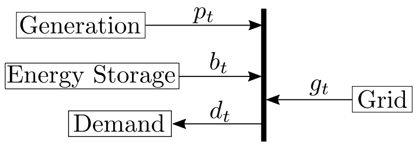

Fig. 1. Block diagram showing positive convention for energy flows. b_t is controlled. b_t and g_t take negative values when energy is transmitted to the battery and grid.

of a receding horizon controller, and its effectiveness at accounting for forecast uncertainty is evaluated for an empirical problem instance using real PV and residential demand data [10];

(ii) A principled (Lloyd-Max Quantization based) approach to selecting forecast scenarios from a distribution of forecast errors is presented, and is evaluated in terms of the performance of a stochastic DP making use of the resulting scenario forecasts;

(iii) The empirical performance of these two approaches, when applied with simple univariate forecasting methods, are demonstrated.

The remainder of this paper is structured as follows: Section II presents the operational optimization approach used, and the two approaches considered for dealing with forecast uncertainty; Section III provides details of the empirical case study considered, and presents the results obtained from simulations; and Section IV draws some conclusions.

II. Method

This section presents a Stochastic Dynamic Programming (SDP) approach to optimally operate an energy storage system, this formulation has been previously presented, [9], and is reproduced here for completeness. A residential user with a rooftop PV system is considered, the sign convention for energy flows is provided in Fig. 1. From Section II-D, onwards, the approaches to accounting for forecast uncertainty, which extend previous work and are a contribution of this paper, are presented.

A. Definitions [units where appropriate]

- Δt — duration of an interval [hours];
- $\mathbb{T}$ — the set of intervals in the horizon, represented by integers $\{0, 1, \ldots, T-1\}$ (there are T intervals; index T is used to refer to the ending state of the final interval);
- d_t^m — forecast demand during interval t in scenario m [kWh];
- p_t^n — forecast generation during interval t in scenario n [kWh];
- $\mathbb{P}(.)$ — marginal probability of outcome $(.)$;
- c_t — cost of buying grid-supplied electricity during interval t [$[1]/kWh];
- r_t — reward for exporting electricity to the grid during interval t [$/kWh];
- η_c, η_d — charge, discharge efficiency of the battery;
- q_t — amount of energy in the battery at the start of interval t (above the lowest allowable amount of energy) [kWh];
- B — usable battery capacity [kWh];
- b_t — decision variable - energy to withdraw from the battery during interval t, (measured at the battery, *i.e.* net of losses when charging, and includes them when discharging) [kWh];
- $b^{\mathrm{m(ax|in)}}$ — maximum and minimum decrease in energy of the battery over an interval (*i.e.* $b_t \in [b^{\min}, b^{\max}]$) [kWh].

[1]Throughout this paper $ refers to Australian Dollars

B. Assumptions

To simplify the presentation of the method, and to keep the problem tractable, the following assumptions are made:

- Self-discharge losses in the battery are negligible;
- Power converter losses are lumped with battery (dis)charging efficiency;
- Battery (dis)charging losses are treated as fixed percentage power losses;
- Realized generation and demand are independent, conditional on forecasts made at $t = 0$;
- Demand and generation are stage-wise independent, conditional on forecasts made at $t = 0$.

Assuming stage-wise independence dramatically simplifies the problem formulation. The independence of realized values, conditional on forecasts made at the beginning of the horizon, is reasonable for early intervals of the horizon (as no information is ignored). Later in the horizon, it is less justified (because if it were possible to re-forecast for later intervals, given realized values of earlier ones, the forecasts would likely change). This increasing approximation error is somewhat mitigated by using this SDP formulation within a receding-horizon controller; any decision made t intervals into the horizon has $t-1$ opportunities for recourse before it is implemented.

C. SDP Formulation

The optimization is formulated as a SDP, with the energy in the battery, q_t, providing a complete description of the system *state*, and the interval within the horizon, t, representing the *stage*. q_t is discretized, therefore the amount of energy transferred in/out of the battery during an interval, b_t, is chosen from a finite set of values.

Denote by $CTG_t(q_t)$, the minimum expected cost-to-go from stage t given that the battery is charged to q_t kWh at the start of interval t. Also, denote by $STC_t(b_t)$ the state-transition-cost for stage t, if control action b_t is chosen (STC also depends on the realized values of generation and demand, but these are omitted for brevity of exposition). The recursive relationship for the minimum cost-to-go is then:

$$CTG_t(q_t) = \min_{b_t}\{\mathbb{E}[STC_t(b_t)] + CTG_{t+1}(q_t - b_t)\} \quad (1)$$

where $\mathbb{E}[.]$ takes the expected value (over possible realizations of PV output, p_t and local demand, d_t). Included in (1) is the state update rule; if one discharges b_t kWh from a battery which contains q_t kWh, it results in a battery containing $(q_t - b_t)$ kWh, one interval later.

The amount of energy provided to the grid when b_t kWh of energy are withdrawn from the battery is represented as $\hat{b}_t$. This can be calculated as:

$$\hat{b}_t = \begin{cases} b_t/\eta_c & b_t < 0 \\ b_t\eta_d & b_t \geq 0 \end{cases} \quad (2)$$

where η_c, $\eta_d \in (0,1]$ are the charging and discharging efficiencies of the battery, respectively.

The expected value of the state transition cost is then:

$$\mathbb{E}[STC_t(b_t)] = \sum_{m=1}^{M}\sum_{n=1}^{N}(\mathbb{P}(d_t^m)\mathbb{P}(p_t^n)$$
$$(c_t[d_t^m - p_t^n - \hat{b}_t]^+ - r_t[d_t^m - p_t^n - \hat{b}_t]^-)) \quad (3)$$

where $[.]^+$ represents taking the positive component, *i.e.* $[x]^+ = \max(0, x)$, and analogously $[x]^- = \max(0, -x)$. In (3) a sum is taken over possible values of the stochastic variables, *i.e.* summing over the $M \cdot N$ possible realizations of demand and generation in interval t; with each weighted by its probability of occurrence. The two terms within the sum are (i) the cost of grid imports, and (ii) the reward for any grid export (at most one of these terms can have a non-zero value).

Satisfaction of system constraints is ensured as follows:

- Energy balance is achieved by defining energy from the grid over interval t as: $g_t = d_t^m - p_t^n - \hat{b}_t$, in (1), and asserting that the grid can supply or accept the required amount of energy;
- Charging and discharging rates are kept feasible by limiting the minimum (most negative) and maximum values of b_t, *i.e.* $b_t \in [b^{\min}, b^{\max}]$;
- Battery capacity constraints are satisfied by additionally limiting the minimum and maximum values of b_t for particular charge states, *i.e.* :

$$\begin{aligned} b_t &\leq \min(q_t, b^{\max}) \\ b_t &\geq \max(q_t - B, b^{\min}) \end{aligned} \quad (4)$$

The optimization of expected cost-to-go (1) is achieved using backwards recursion. The minimal cost-to-go from the end of the last interval in a horizon, $CTG_T(q_T)$, is zero for all ending states, q_T (there can be no further cost as the horizon is complete). For each preceding interval, $t = \{T-1, \ldots, 0\}$, and each possible charge-level at the start of that interval, q_t, the feasible discharge decision, b_t, that minimizes (1) is found by exhaustively searching the finite set of possible values. This process is repeated, moving backwards through the horizon, until the minimal cost-to-go from the first interval of the horizon $CTG_0(q_0)$, is found (the starting battery charge-level, q_0, is known from the simulation).

This SDP is solved for each horizon as a receding horizon controller; once solved, the first decision, b_0, is applied in an on-line simulation, and a new solution is computed one interval later, with the horizon and forecasts moved forward one interval.

D. Approaches to Accounting for Forecast Uncertainty

In addition to solving the optimization problem using a receding horizon (enabling recourse of decisions at intervals later in the horizon), the present work considers two approaches to explicitly dealing with the inevitable uncertainty associated with forecasts:

1) Cost Discounting: A common approach in receding horizon operational optimization is to apply a discount factor to costs/rewards which are realized at later stages of the horizon. In the present study discounting is applied to account for forecast uncertainty. It is noted that there are other uses for discount factors (*e.g.* accounting for time-value-of-money or model approximation errors). With real forecasts there are two reasons a discount-factor based approach may be effective; the first is that forecasts are generally less accurate further into the future, the second is that the cumulative impact of forecast errors will mean the costs and benefits further into the horizon are increasingly uncertain (even if the uncertainty of forecasts themselves are uniform across the horizon).

To include cost discounting the recursive cost equation is modified to:

$$CTG_t(q_t) = \min_{b_t}\{\mathbb{E}[STC_t(b_t)] + \beta \cdot CTG_{t+1}(q_t - b_t)\} \quad (5)$$

Where β is the discount factor. This discounts the cost from interval t by β^t, where $t \in \{0, 1, \ldots, T-1\}$. The parameter β can be chosen for a particular problem instance based on, for example, simulation on a historical data-set.

2) Scenario Forecasts: The second approach considered to address forecast uncertainty is to look at a number of scenarios of demand and PV output over the horizon (*i.e.* $M, N > 1$ in (1)). To use the proposed SDP formulation, it is necessary to express the probability distribution of stochastic variables as a number of discrete scenarios, and it is not computationally tractable to consider a large number of scenarios.

Given that it is necessary to consider only a small number of forecast scenarios[2], N, it is important that they are chosen effectively, in order to best represent the continuous probability distribution of possible outcomes.

[2]For simplicity we consider an equal number of scenarios of PV and demand, N (resulting in N^2 independent scenarios of PV and demand).

To achieve this, Lloyd-Max-Quantization is applied [11], [12], which is a principled approach that minimizes the expected squared-error of discretization. It is described briefly here for the reader's convenience:

For a given continuous random variable X with probability density function $f(x)$ (*i.e.* $\mathbb{P}(a \leq X \leq b) = \int_a^b f(x)dx$), which is to be represented as N discrete scenarios (*i.e.* with a discretized probability-mass function), we seek the $N-1$ decision boundaries[3] $\{b_k\}_{k=1}^{N-1}$ and N reconstruction levels $\{y_k\}_{k=1}^{N}$ that minimize the expected squared-error of discretization, known in the signal processing literature as Distortion:

$$D = \mathbb{E}[(x - Q(x))^2] = \int_{-\infty}^{\infty} (x - Q(x))^2 f(x) dx = \sum_{k=1}^{N} \int_{b_{k-1}}^{b_k} (x - y_k)^2 f(x) dx \quad (6)$$

Where $Q(x)$ represents the quantization of the random variable X, *i.e.* :

$$Q(x) = y_k, \quad x \in [b_{k-1}, b_k), \quad \forall k \in \{1, \dots, N\} \quad (7)$$

Finding the parameter values $\{b_k, y_k\}$ which minimize Distortion can be done exactly, for any known continuous probability density function $f(x)$, and any chosen number of quantization levels, N, by making use of the Lloyd-Max algorithm [11], [12]. Algorithm 1 presents the Lloyd-Max algorithm applied to a univariate distribution, but it can be extended to multivariate distributions to give the Linde-Buzo-Gray algorithm [13]. Algorithm 1 minimizes Distortion, (6), by iteratively solving a pair of partial differential equations, each of which finds the parameters b_k and y_k (respectively) which minimizes Distortion, whilst keeping the other set of parameters constant.

E. Univariate Forecasting and Forecast Errors

In this paper only simple univariate forecasting methods are considered. Univariate forecasts are of particular interest to the operation of distributed energy systems, because they can be produced in the absence of additional communication infrastructure (as they rely on locally available information only). Two forecasts are considered:

1) Naive (Daily) Periodic (NP): forecasts assume that the demand/PV output during any interval is identical to that from 24-hours ago. This crudely captures the major (daily) seasonal period, and provides a naive benchmark.

[3]Implicit boundaries of $b_0 = -\infty$ and $b_N = \infty$ are included, so there are $N+1$ boundaries, but only $N-1$ of them need to be chosen.

Algorithm 1 Selecting the Scenario Values & Probabilities of Occurrence

Require: $N, \epsilon, f(x)$

$y_k \leftarrow y_{k0} \quad \forall k \in \{1, \dots, N\}$ ▷ Initialize[4] y_k
$b_0 \leftarrow -\infty, b_N \leftarrow \infty$
$\Delta_{\text{rel.}} \leftarrow \infty$
while $\Delta_{\text{rel.}} > \epsilon$ **do**
    $b_k \leftarrow 0.5(y_k + y_{k+1}) \quad \forall k \in \{1, \dots, N-1\}$
    $y_k^{\text{prev}} \leftarrow y_k \quad \forall k \in \{1, \dots, N\}$
    $y_k \leftarrow \int_{b_{k-1}}^{b_k} x f(x) dx / \int_{b_{k-1}}^{b_k} f(x) dx \quad \forall k$
    $\Delta y \leftarrow \sum_{k=1}^{N} (y_k - y_k^{\text{prev}})^2$
    $\Delta_{\text{rel.}} \leftarrow \Delta y / \sum_{k=1}^{N} (y_k^{\text{prev}})^2$
$\mathbb{P}(y_k) \leftarrow \int_{b_{k-1}}^{b_k} f(x) dx \quad \forall k \in \{1, \dots, N\}$
return $\{y_k, \mathbb{P}(y_k)\} \quad \forall k$

2) Multiple Linear Regression (MLR): forecasts assume the demand/PV output over the horizon is some linear combination of the previous D horizons of realized values[5]. The forecast is:

$$\boldsymbol{F} = \boldsymbol{X\beta} \quad (8)$$

Where $\boldsymbol{F} \in \mathbb{R}^{1 \times T}$ is a row vector of point forecast values for the next T intervals, $\boldsymbol{X} \in \mathbb{R}^{1 \times TD}$ is a row vector of realized values for the previous TD intervals, and $\boldsymbol{\beta} \in \mathbb{R}^{TD \times T}$ is a matrix of parameters selected for a particular customer using linear regression on a training dataset. For the data considered in this paper $D = 7$ days provided reasonable performance for both PV and demand forecasting (on an unseen validation dataset), and was used in all results presented.

3) Perfect Foresight (PF): forecasts with zero error are used to provide a non-tight upper bound on the performance that could be achieved with better forecasts (*i.e.* the performance of the approach under perfect information).

4) Forecast Errors: To produce scenario forecasts as described in Section II-D2, it is necessary to have a probability density function of the demand and generation in each interval within the horizon. To generate this, a statistical model is assumed for the realized PV output or demand over the horizon, $\boldsymbol{Y}$, as:

$$\boldsymbol{Y} = \boldsymbol{F} + \boldsymbol{\epsilon} \quad (9)$$

Where $\boldsymbol{F}$ is the output from a point forecast model, and $\boldsymbol{\epsilon}$ is a vector of random variables representing the

[4]y_k could be initialized randomly. For the Normal distribution N linearly spaced values over $[\mu - 2\sigma, \mu + 2\sigma]$ offered faster convergence.

[5]In general MLR can include additional regressors such as forecast temperature, dummy variables to encode day of the week, holidays, etc. For simplicity the univariate case is considered.

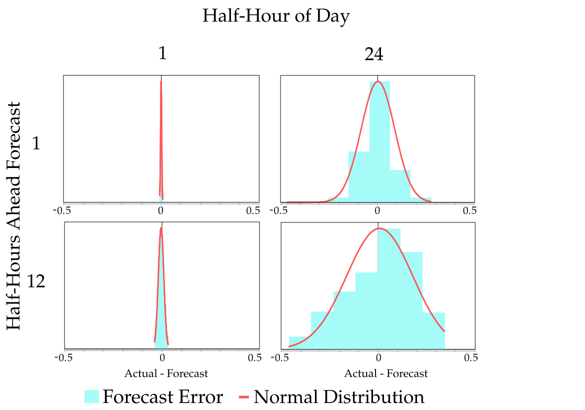

Fig. 2. Forecast Error Distribution for MLR PV Forecasts for a Typical Customer, for 4 of the 48×48 Error Groups. Indicates that assuming a Normal distribution for these errors is a reasonable approximation. During the middle of the day (half-hour No. 24), PV forecast errors are greater, and they are slightly reduced for nearer-term forecasting.

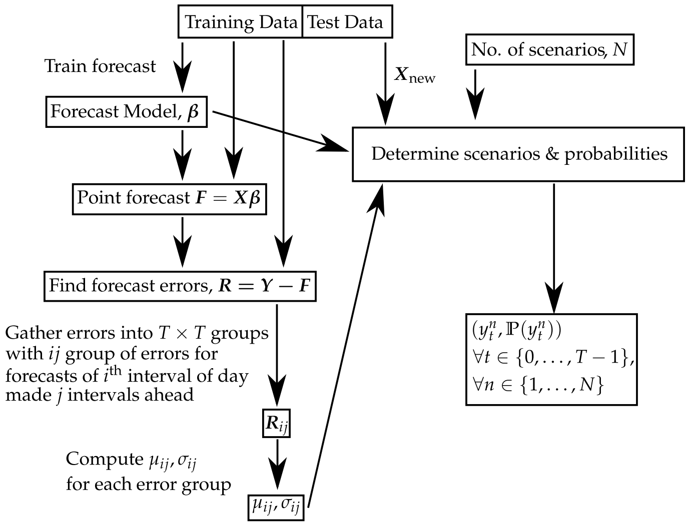

Fig. 3. Producing MLR Scenario Forecasts. X_{new} is a feature vector for a forecast horizon not included in the training data.

forecast errors in each step of the horizon. To improve the resolution of scenario forecasts, errors ϵ are divided into groups, and a Normal distribution is fitted to each group. The forecast error groups are defined as follows:

For NP forecasts, forecast errors are divided into groups based on the interval-of-the-day for which a forecast is made (so for half-hour intervals we have 48 groups of errors). Low (zero) forecast errors can be expected for the PV output during the hours of darkness.

For MLR forecasts, forecast errors are further subdivided based on the number of intervals in advance the forecast is made. Here nearer-term forecasts of PV and demand might be expected to have lower forecast errors.

Once divided into their groups, a Normal distribution is assumed for the forecast errors, ϵ. The standard deviation for each group of errors is based on the performance of the chosen forecast model on the training data[6].

Fig. 2 gives a few typical forecast error histograms from MLR forecasting of PV output (on the training data) and the Normal distribution by which they are represented when producing scenarios. These results indicate that assuming a Normal distribution for forecast errors from the point forecasting models is reasonable.

The overall process of producing MLR scenario forecasts for the SDP is shown schematically in Fig. 3.

III. Case Study & Results

Table I summarizes the features of the empirical simulations used to assess the two approaches to addressing forecast uncertainty. The results presented are in terms of the cost-savings delivered to the residential customer, compared to not having a battery available (*i.e.* immediately exporting any excess PV generation, and importing to meet excess local demand). Two years of half-hourly

TABLE I
Baseline Simulation Parameters

Parameter	Symbol	Value	Units
Interval Length	Δt	0.5	hours
Horizon Length	T	48	intervals
Import Price	c_t	0.4 (7AM-10PM) 0.2 (otherwise)	\$/kWh
Export Price	r_t	0.05	\$/kWh
Battery Capacity	B	2	kWh
(Dis)Charging efficiency	η_c, η_d	0.94	[]
Maximum charge current	$I_{ch,\max} = -b^{\min}/(B\Delta t)$	1.0	C
Maximum discharge current	$I_{d,\max} = b^{\max}/(B\Delta t)$	2.0	C

demand and PV output are taken for 16 customers for 2011-2013 from the Ausgrid 'Solar Home Electricity Data' [10]. Peak demands for the customers was between 4 and 9kW, and peak PV output between 1 and 2.5kW. The first year is used for training the MLR forecast model, and to fit the forecast error Normal distributions for the MLR and NP models. The second year is used to simulate the performance of the battery controlled using the SDP method outlined above, provided with various forecasts.

Fig. 4 shows the standard deviation of PV and demand forecast errors for an example customer, after an MLR forecast model has been fitted to the training data. As might be expected forecast errors are highest at the times of day when demand and PV output are generally greater, and there are some slightly reduced forecast errors for near-term forecasting.

Fig. 5 shows the cost saved by a battery operated using SDP provided with point ($N = 1$) versions of the NP, MLR, and the PF forecasts, as the discount factor, β, is varied. For the PF forecast results are as expected, reduc-

[6] For the MLR forecast this is expected to under-estimate forecast errors, when the models are applied to unseen test data

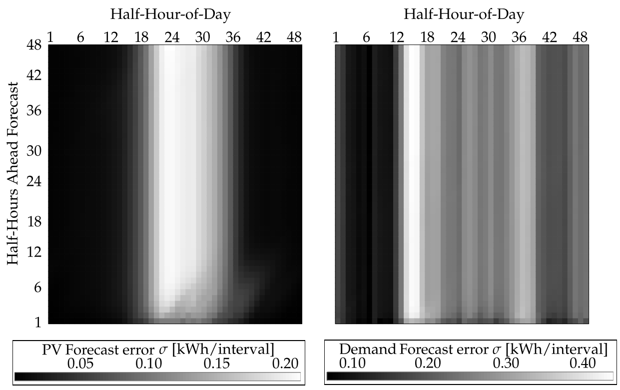

Fig. 4. Standard deviation of PV (Left) and Demand (Right) forecast errors from MLR forecasting, divided into 48×48 groups, based on half-hour interval of the day (horizontal axis), and No. of intervals ahead forecast is made (vertical axis), for a typical customer. There is a clear pattern of increased forecast errors during times of high PV output (8AM-5PM). There are also reduced forecast errors when forecasting in the near-term (bottom of plot). Similar patterns can be seen in the Demand forecast errors, but the benefits of forecasting in the near-term last only an hour, as a result of the volatility of demand for a single household.

ing β from unity causes a loss in performance, because the forecasts have zero error. A similar pattern can be seen for the MLR and NP forecasts, indicating that cost-discounting is an ineffective approach to accounting for forecast uncertainty in this application. This is believed to be the result of two issues; firstly forecast errors are only weakly correlated with the distance into the horizon that forecasts are made for (as can be seen by the relative uniformity of Fig. 4 along the vertical axis), reducing the validity of the cost-discounting approach to dealing with forecast errors. The second issue is that using a discount factor $\beta < 1$ increases the relative importance of costs associated with the first few intervals; and these intervals have the least robustness to forecast errors from receding horizon control. Receding horizon control provides an opportunity for recourse of intervals further into the horizon, making the costs of those intervals less sensitive to forecast errors, whereas solving the optimization with $\beta < 1$, in effect, assumes intervals further into a horizon are more sensitive to forecast errors.

The performance improvement achieved by considering an increasing number of forecast scenarios is shown in Fig. 6. For the MLR forecast increasing the number of scenarios monotonically increases the average performance, and with $N = 7$ scenarios the battery offers 8% greater cost savings per year, on average. For the NP forecast the pattern is less clear, and cases with an even number of scenarios tend to offer reduced performance on average. This is probably because when N is even the expected case is not included as a scenario; for example if we know the PV output during a particular interval is Normally distributed with mean μ and standard deviation σ, the forecast values for $N = 2$ scenarios are

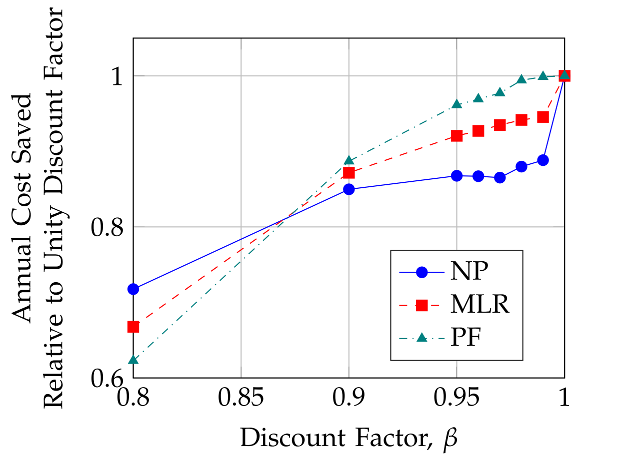

Fig. 5. Value of 2kWh Battery Operated Using SDP with Various Forecasting Methods, over a range of Discount Factors, β. Line shows the value relative to the value achieved at a unity discount factor $\beta = 1$, averaged over the 16 customers.

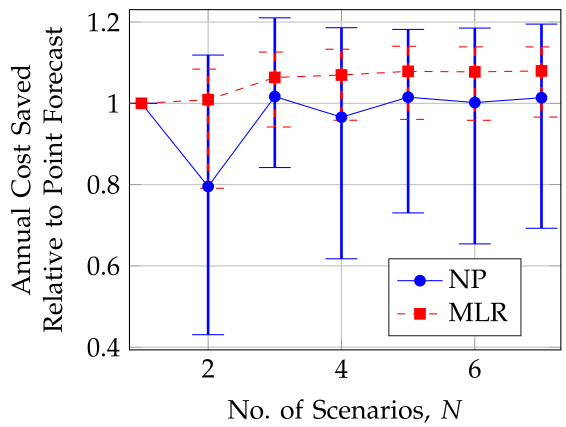

Fig. 6. Value of 2kWh Battery Operated Using SDP with an Increasing Number of Scenarios, using NP and MLR Forecasts, each shown relative to performance of a point forecast of the same type. Line shows average over 16 customers, whiskers show the min/max cases.

$\mu \pm 0.798\sigma$, so do not include the expected PV output, μ. For $N = 7$ scenarios the NP forecast performs just 1% better on average, with a significant spread between the 16 customers (with some achieving worse performance than when using a NP point forecast).

Fig. 7 shows the annual battery cost saving based on a 1-year simulation, using 5 different forecasting methods (and a unity discount factor). MLR point forecasting offers a 35% increase, on average, in cost-savings over NP point forecasting, and considering $N = 7$ scenarios around the MLR point forecasts offers a further 8% improvement. If the SDP were given access to a perfect foresight forecast it can perform substantially better again, but this is a non-tight upper bound on performance as any real forecast will have some non-zero error. The cost savings from a 2kWh (usable capacity) battery, under the tariff as-

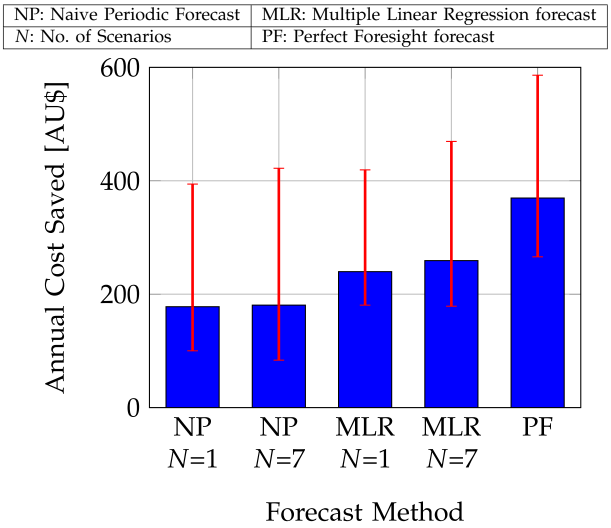

Fig. 7. Value of 2kWh Battery Operated Using SDP with Various Forecasting Methods. Bars show the mean for 16 customers considered, the whiskers show the min/max cases.

sumptions made, are 180-470$/year if MLR forecasts with $N=7$ scenarios are used. If the battery lasts 5-years this corresponds to a lifetime cost-saving of 900-2,350$. At the time of writing commercially available one-off residential battery systems cost approximately 1,000$/kWh (nominal capacity). Therefore some of the significant battery cost reductions which have been forecast, need to be realized before batteries are a cost-effective investment for most Australian home-owners, unless value streams beyond the two considered in the present work are exploited.

IV. Conclusion

Two approaches to accounting for forecast uncertainty when using a receding horizon controller, with dynamic programming for optimizing decisions over a forecast/control horizon, have been presented.

The first approach applies a discount factor across the forecast/control horizon. A below unity discount factor ($\beta < 1.0$), which reduces the relative importance of costs further into the horizon, might be expected to improve performance, as those parts of the horizon are less certain due to forecasting errors. However, this was not found to be the case, and this is primarily the result of receding horizon control providing increasing opportunities for recourse further into the horizon, making costs associated with events further into the horizon, counter-intuitively, less sensitive to forecast errors.

The second approach converts point-forecast models into probabilistic (scenario) forecast models based on Lloyd-Max quantization of the point forecast error distributions; and uses these forecasts to optimize control actions via stochastic dynamic programming. This approach increased the annual cost-savings of a 2kWh battery by 8%, on average, compared to a similar approach which considered point forecasts only.

Acknowledgment

This work was supported by Melbourne International Research and Fee Remission Scholarships, and an Australia-Indonesia Center grant.

References

[1] P. Mirowski, S. Chen, T. K. Ho, and C.-N. Yu, "Demand forecasting in smart grids," *Bell Labs Technical Journal*, vol. 18, no. 4, pp. 135–158, 2014.

[2] A. Tuohy, J. Zack, S. E. Haupt, J. Sharp, M. Ahlstrom, S. Dise, E. Grimit, C. Mohrlen, M. Lange, M. G. Casado, J. Black, M. Marquis, and C. Collier, "Solar forecasting: Methods, challenges, and performance," *IEEE Power and Energy Magazine*, vol. 13, no. 6, pp. 50–59, Nov 2015.

[3] N. V. Sahinidis, "Optimization under uncertainty: state-of-the-art and opportunities," *Computers & Chemical Engineering*, vol. 28, no. 6-7, pp. 971–983, 2004, {FOCAPO} 2003 Special issue. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0098135403002369

[4] A. Oudalov, R. Cherkaoui, and A. Beguin, "Sizing and optimal operation of battery energy storage system for peak shaving application," in *Power Tech, 2007 IEEE Lausanne*, July 2007, pp. 621–625.

[5] V. Muenzel, I. Mareels, J. de Hoog, A. Vishwanath, S. Kalyanaraman, and A. Gort, "PV generation and demand mismatch: Evaluating the potential of residential storage," *2015 IEEE Power & Energy Society Innovative Smart Grid Technologies Conference (ISGT)*, p. 1, 2015.

[6] J. Qin, R. Sevlian, D. Varodayan, and R. Rajagopal, "Optimal electric energy storage operation," in *Power and Energy Society General Meeting, 2012 IEEE*, July 2012, pp. 1–6.

[7] X. Xi, R. Sioshansi, and V. Marano, "A stochastic dynamic programming model for co-optimization of distributed energy storage," *Energy Systems*, vol. 5, no. 3, pp. 475–505, 2014.

[8] L. Hannah and D. B. Dunson, "Approximate dynamic programming for storage problems," in *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, 2011, pp. 337–344.

[9] K. Abdulla, J. D. Hoog, V. Muenzel, F. Suits, K. Steer, A. Wirth, and S. Halgamuge, "Optimal operation of energy storage systems considering forecasts and battery degradation," *IEEE Transactions on Smart Grid*, vol. PP, no. 99, pp. 1–1, 2016. [Online]. Available: http://dx.doi.org/10.1109/TSG.2016.2606490

[10] Ausgrid solar home electricity data, v.2. (Accessed: 15th Sept. 2015). [Online]. Available: http://www.ausgrid.com.au/Common/About-us/Corporate-information/Data-to-share/Data-to-share/Solar-household-data.aspx

[11] S. Lloyd, "Least squares quantization in pcm," *IEEE Transactions on Information Theory*, vol. 28, no. 2, pp. 129–137, Mar 1982.

[12] J. Max, "Quantizing for minimum distortion," *IRE Transactions on Information Theory*, vol. 6, no. 1, pp. 7–12, March 1960.

[13] Y. Linde, A. Buzo, and R. Gray, "An algorithm for vector quantizer design," *IEEE Transactions on Communications*, vol. 28, no. 1, pp. 84–95, Jan 1980.