Optimal Control of Residential Energy Storage Under Price Fluctuations

Apoorv Srivastava and Dhvaneel Visaria

{apoorv1, dhvaneel}@stanford.edu

June 7, 2023

Abstract

The increasing demand for energy, coupled with the growing concerns over environmental sustainability, has led to a significant shift in the global energy landscape. As renewable energy sources, such as solar and wind, gain prominence, the variability and intermittency of these sources pose challenges to grid stability. In addition, demand trends reveal large difference in power required during the evening compared to the day, which further exacerbates the grid stability. To address these issues and optimize energy utilization, the integration of energy storage systems in residential settings has emerged as a promising solution. Residential energy storage systems (RESS) enable homeowners to store excess energy during periods of low demand and release it during peak hours or when renewable energy generation is low. By effectively managing the energy flow within these systems, homeowners can reduce their reliance on the grid during peak hours, optimize their energy consumption, and potentially save costs.

1 Introduction

One crucial factor in residential energy storage management is the price fluctuations of electricity. In many electricity markets, electricity costs vary throughout the day due to factors such as demand patterns, generation capacity, unforeseen circumstances, and market conditions. These price fluctuations provide an opportunity for homeowners to strategically schedule the charging and discharging of their energy storage systems to maximize cost savings. Optimal control strategies that consider price fluctuations can significantly enhance residential energy storage's economic viability and efficiency. This project aims to explore and develop an optimal control framework for residential energy storage systems under price fluctuations, with the objective of minimizing electricity costs and variability of grid demand. By investigating different hyperparameters of an optimal control algorithm and evaluating its performance, this study seeks to contribute to the development of more intelligent and economically viable residential energy storage strategies.

2 Related Work

The pressing importance of the grid optimization problem has led to the exploration of various optimization techniques and algorithms, of which Model Predictive Control (MPC) emerges as one of the most promising approaches. MPC [1] utilizes a mathematical model of the energy storage system, along with forecasts of electricity prices and household energy demand, to formulate an optimization problem. The MPC algorithm generates a sequence of control actions over a specific prediction horizon by considering constraints such as the state of charge (SoC) limits, maximum charging/discharging rates, and grid connection limits. The optimization problem is solved iteratively, updating the control actions at each time step after running an open-loop control problem for a finite horizon.

Optimized strategies for power usage allow homeowners to leverage price fluctuations and dynamically adjust the charging and discharging profiles of the energy storage system, aiming to minimize electricity costs without compromising energy availability and, at the same time, reduce the variance in power drawn from the grid, thus enhancing grid stability. Alternatively, the residential energy storage problem has also been addressed through Dynamic Programming (DP) [2], which involves decomposing the original optimization problem into smaller sub-problems and solving them recursively. The optimal solutions for the sub-problems are combined to obtain a solution for the original optimization problem. In finite-horizon problems with or without stochasticity, DP-based techniques are known to be effective.

Besides the optimal control techniques, learning-based control techniques have proven to be very effective as well. Online feedback optimization [3] is one such technique implemented to use measurements from the grid as feedback to iteratively change the control inputs until they converge to the solution of the optimization problem. This method is computationally light, robust to model mismatch, can utilize a grid to its full capacity, and needs very limited model information. The newly realized power of machine learning and artificial intelligence, coupled with the abundance of data in the digital age, has unlocked new avenues for the optimal control of residential energy storage systems. Reinforcement Learning (RL) algorithms [4], such as deep Q-learning and policy gradient methods, are shown to learn optimal control policies through interactions with the environment without requiring an explicit mathematical model. These algorithms can adapt to changing electricity price patterns and household energy demand profiles, continuously improving the control strategies.

3 Problem Statement

In this work, we consider n households, each anchored to a rechargeable battery and a home solar electric system and connected to a grid. Given demand patterns for each home along with the trend of generated power from the solar electric system, we aim to find an optimal charge-discharge policy for the batteries which minimizes the total cost of electric power drawn from the grid and at the same time minimizes the large power load on the grid supply during peak hours.

We evaluate grid operation and the cost minimization problem over a week, with an hourly discretization in time, with the start of the hour indicated by $t^k, k = 0, \ldots, N$. All power supplies and demands are assumed constant over any given time step $\Delta t = 1$ hour. For i^{th} household during the time t_{k-1} to t_k, the power demand D_k^i is fulfilled using supplies from the grid X_k^i, the solar electric system S_k^i and the battery Y_k^i. It is important to note that while supplies from the grid and solar electric system are unidirectional, power can flow into the battery to charge it, i.e., Y_k^i can be negative, as shown in fig. 1. A power sink Z_k^i is added to account for any residual left after fulfilling the household demand and charging the battery and is needed to ensure the safety of the system, for instance, in the case when the battery is fully charged, and power generated by the solar electric system is higher than the demand. The resulting constraints can then be described

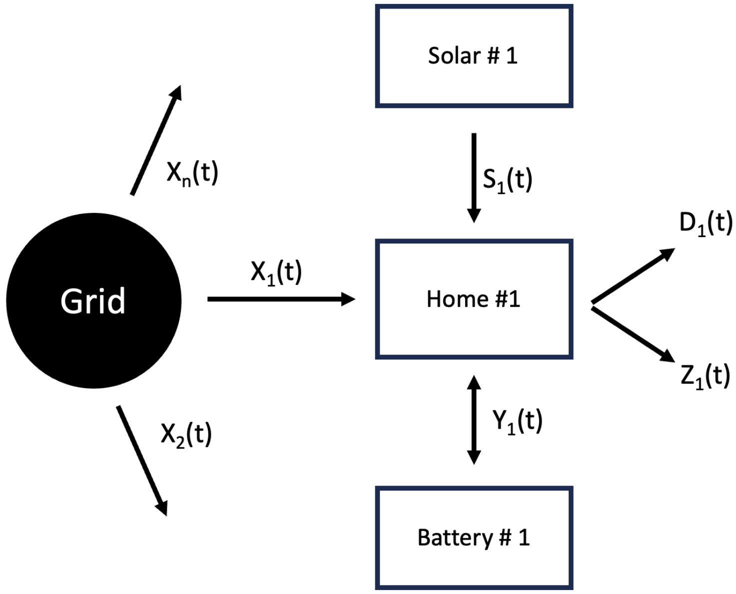

Figure 1: Graphical representation of the model

as, $\forall i = 1, \ldots, n$ and $\forall k = 1, \ldots, N$, we require that

$$D_k^i = X_k^i + S_k^i + Y_k^i - Z_k^i, \tag{1}$$

$$X_k^i \geq 0, \quad S_k^i \geq 0, \quad Z_k^i \geq 0. \tag{2}$$

The battery capacity and energy (or buffer) level at the start of k^{th} hour, for i^{th} household, is indicated by $\bar{B}^i$ and B_k^i, respectively. Each battery is characterized by buffer limits $B_{\min}^i$ and $B_{\max}^i$, maximum charge/discharge rate $\bar{Y}^i$, and charging and discharging efficiency η_c and η_d, respectively. The inclusion of charging and discharging efficiency in the model highlights the inefficiencies involved and establishes correspondence with available technology, while limits on buffer levels and maximum charge/discharge rate ensure the longevity of the RESS. Following this, we get the battery energy level evolution equation along with the associated constraints $\forall i = 1, \ldots, n$ and $\forall k = 1, \ldots, N$ as

$$B_k^i = B_{k-1}^i - f(Y_k^i), \quad f(Y) = \max\{\eta_c Y,\ Y/\eta_d\}, \tag{3}$$

$$B_{\min}^i \leq B_k^i \leq B_{\max}^i, \tag{4}$$

$$-\bar{Y}^i \leq Y_k^i \leq \bar{Y}^i. \tag{5}$$

We look at minimizing the total cost and the peak load on the grid to ensure grid stability. With total power drawn from the grid between time t_{k-1} and t_k described by $Q_k = \Sigma_{i=1}^n X_k^i$ and given unit price of power p_k over the time step, aggregated in vectors $\mathbf{Q} = [Q_1, \ldots, Q_N]$ and $\mathbf{p} = [p_1, \ldots, p_N]$, total power cost is given by $\mathbf{p}^\top \mathbf{Q}$. Consequently, the objective function for the optimization problem can then be expressed as

$$\text{minimize} \ \ \alpha_1 |\mathbf{Q}|_\infty + \alpha_2 \mathbf{p}^\top \mathbf{Q}, \tag{6}$$

where α_1 and α_2 are the weighing factors for objectives corresponding to peak demand and total cost, respectively.

Qualitatively, the battery is expected to draw power from the grid during low-demand periods and consequently use the stored energy during the peak demand hours when the per-unit cost is high. Since a fraction of the power demand during the peak hour is supplied by the battery, this will also have a balancing effect on the grid. In other words, it will minimize the difference between the power supplied by the grid during peak and nadir demand hours.

4 Method

The hourly electricity prices for a week, required as input for the problem, are obtained using the normalized prices identified using [5] and the average electricity price in the United States. We assume that the normalized demand follows a similar trend by following a supply-demand relation, i.e., the per-unit price goes up because demand increases. This, along with the average American home consumption of 30 kWh, as reported by the Energy Information Administration, is used to obtain the total hourly base demand input for a week. It is important to note that the demand patterns during the weekend are significantly different from those during the week, as shown in fig. 2. Hourly input data corresponding to normalized solar power generated during a week is obtained from a report published by the California ISO [1]. To account for the inherent variability associated with demand and generated solar power, corresponding inputs for each household are corrupted with independent random Gaussian noise.

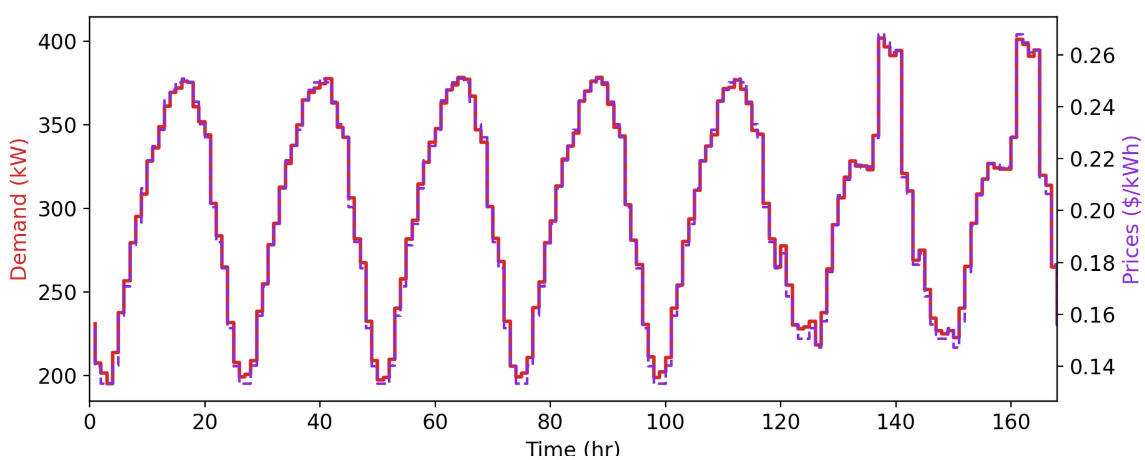

Figure 2: Demand, D(t) and electricity price, p(t) data for a week with a similar trend depicting positive correlation and different trends between weekdays and weekends

The optimization problem with objective function given by (6) subjected to constraints described by (1)-(5) is convex. However, the constraint describing the evolution of battery energy levels, given by (3), is difficult to handle under the modeling restrictions associated with any optimization algorithm. Considering this, we use Model Predictive Control (MPC), in addition to convex

[1]What the duck curve tells us about managing a green grid, California ISO

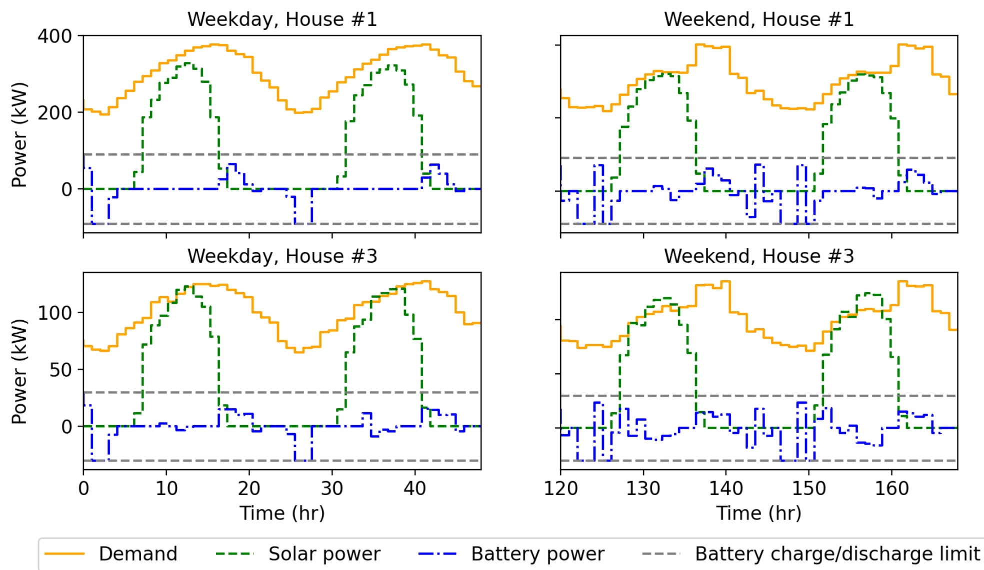

Figure 3: Demand, solar power, and optimal battery power for two weekdays (left column) and weekends (right column) for two different homes

programming, to obtain an optimal energy storage policy.

The nonlinearity in (3) is isolated in $f(Y)$, which admits a bi-linear form. We aim at finding a linear approximation $\hat{f}(Y)$ for $f(Y)$, which ensures that buffer and charge/discharge rate limits of the battery described by (4) and (5) are not violated. This is important numerically to guarantee the feasibility of the optimization problem and practically, to ensure the safety and longevity of the battery. Using knowledge of the model and battery dynamics, it is identified that $0 \leq \hat{f}(Y) \leq f(Y)$ for $Y \geq 0$ and $0 \geq \hat{f}(Y) \geq f(Y)$ for $Y < 0$ to ensure battery energy level and charge/discharge rate constraints. Consequently, approximation $\hat{f}(Y)$ is given by

$$\hat{f}(Y) = aY, \quad a = \min\{\eta_c, 1/\eta_d\}. \tag{7}$$

Replacing $f(Y)$ by $\hat{f}(Y)$ in (3) results in a straightforward convex problem with convex objective function and affine constraints, which is solved using CVXPY[6][7]. Under the MPC framework, true battery energy levels after one time step are estimated according to (3) using known initial energy level and the optimal values of the power flowing from the battery for the first step Y_0^{opt} obtained from the convex program.

5 Experiments

The results discussed below are generated using $n = 5$ households. Varying sizes of the houses connected to the grid are accounted for by randomly choosing (with replacement) n number of battery capacities $\bar{B}^i$ from a list of standard battery capacities (100kWh, 300kWh, and 500kWh). Battery properties are assigned based on battery capacity as $B^i_{\text{min}} = 0.2\bar{B}^i$, $B^i_{\text{max}} = 0.8\bar{B}^i$, and $\bar{Y}^i = 0.3\bar{B}^i/\Delta t$. Power demand and solar power generated are also scaled in proportion to the battery capacities. We synthesize the results for a week using an equally weighted objective function, i.e., $\alpha_1 = \alpha_2 = 1$, and MPC with time horizon $N_{\text{MPC}} = 24$ hours.

The power supplied from the battery according to the optimal policy under the aforementioned conditions, as the demand pattern unfolds, is shown in fig. 3. As we can see, the battery is charged and discharged to offset the power required from the grid during peak hours. Initially, the battery is discharged to provide for the demand when solar power is unavailable. In the afternoon, as the solar power becomes accessible and is in excess of the required demand, the battery is charged, as can be observed clearly for House #3 during the weekend (right column) between hours 125 and 135. It is important to note that the battery is being charged during peak demand periods, but this energy is coming from the solar panels, and not the grid, and hence is effectively saving cost. Through the peak demand period in the evening, the battery is discharged to accommodate the high power requirements.

Changes in the energy levels of the battery as the power is drawn from or supplied to the battery is shown in fig. 4. Horizontal dashed lines show energy level limits and charge/discharge rate limits. The battery energy levels often hit upper and lower limits, indicating that the battery is being effectively used, and energy level transitions and charge/discharge rate show that the battery dynamics and charge/discharge rate constraints are respected.

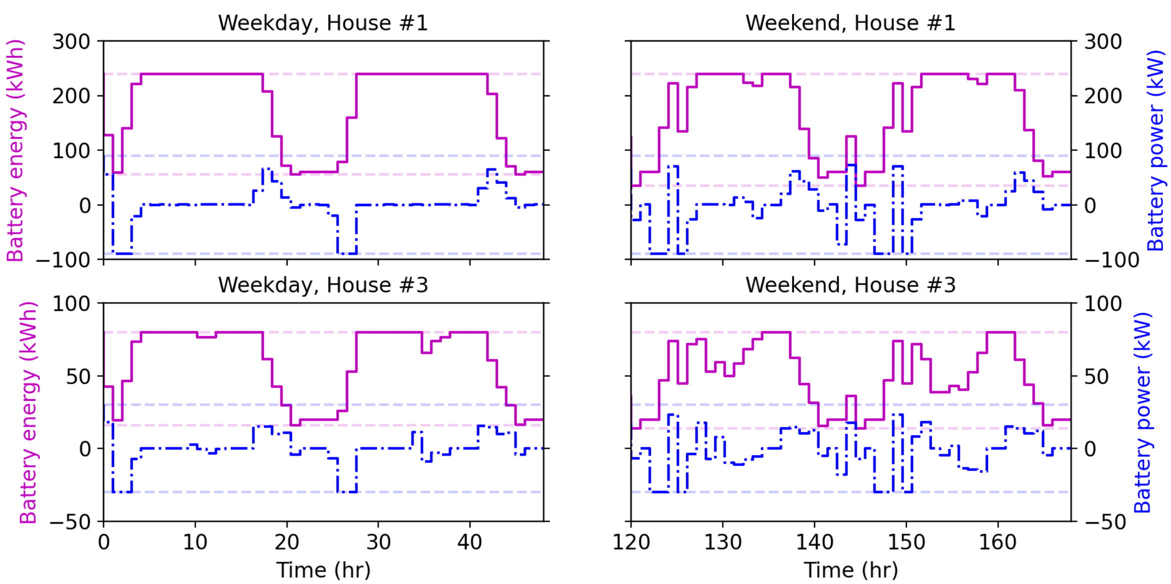

Figure 4: Optimal battery energy level and optimal battery power for two weekdays (left column) and weekend (right column) for two different homes with buffer limits and battery charge/discharge rate limits

We investigate the differences in results from the optimal policies focusing on minimizing the power cost with those more concerned with grid stability by varying weights α_1 and α_2 in the objective function. Three policies, namely cost-focused policy ($\alpha_1 = 0$, $\alpha_2 = 1$), grid-focused policy ($\alpha_1 = 1$, $\alpha_2 = 0$), and equally weighted policy ($\alpha_1 = 1$, $\alpha_2 = 1$), are evaluated and total grid supply under different policies are reported in fig. 5. Under the cost-focused policy, peak grid supply is much higher compared to the other two policies. On the other hand, under the grid-focused policy, the maximum power drawn from the grid is much lower, but significantly more power is drawn during low-demand hours. Equally weighted policy embodies the best of the two extreme policies, i.e., the peak load is curtailed while lowering the price.

Changes in policy under a higher number of households n and longer time frame N were also inspected. It was observed that the optimal policy is invariant to changes in these parameters. The independence of the policy under n and N can be argued under the assumption that the grid is an infinite source of power and total time doesn't affect optimal output coming from MPC, looking at only N_{MPC} steps ahead, respectively. With respect to the MPC time horizon N_{MPC}, the optimal control policy is more responsive, with finer fluctuations, with larger N_{MPC} since the optimization problem looks further into the future and has access to more information. However, this incremental accuracy comes at the cost of significantly higher computational effort.

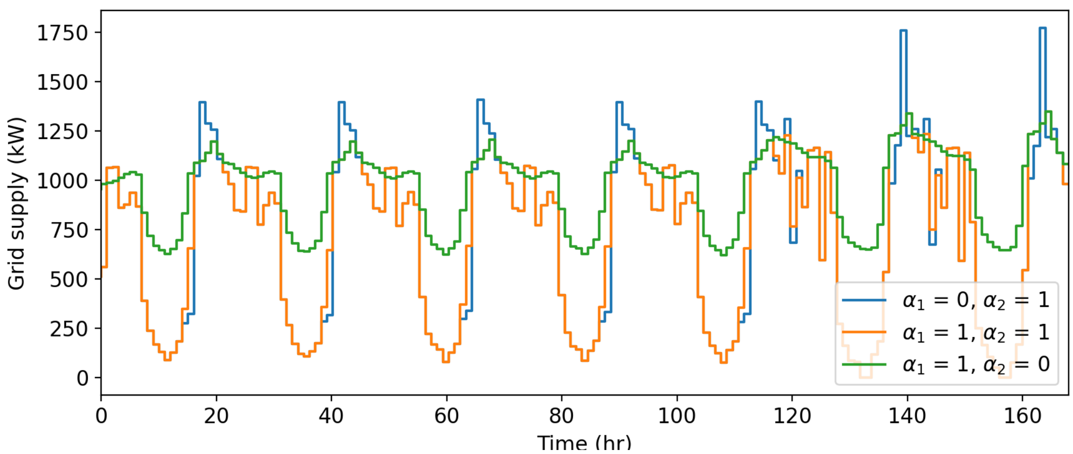

Figure 5: Total grid supply under the cost-focused, grid-focused, and equally weighted policy.

6 Conclusions and Future Work

In summary, this study aims to enhance previous research conducted on the problem of residential energy storage, presented in [8]. This work considers an electric power grid connected to multiple houses equipped with solar electric systems and rechargeable batteries and focuses on optimizing the charge-discharge policy of rechargeable batteries under uncertainties and variability in the demand and solar power generated. The discretized problem is solved using Model Predictive Control (MPC) in conjunction with convex programming. A linear approximation of the battery dynamics is utilized to ensure feasibility and safety constraints. It is important to note that sequential convex programming is ineffective since the power flowing from the battery is not continuous and changes from one step to the next, although bounded, can be large.

Experimental results demonstrate the effectiveness of the optimized battery power policy in offsetting power from the grid during peak hours. The battery is charged during low-demand periods when solar power is abundant and discharged during peak-demand periods. The battery energy levels adhere to their limits, and the charge/discharge rates remain within the specified bounds. Different weighting factors in the objective function are explored, highlighting the trade-off between minimizing power cost and ensuring grid stability. The study also examines the impact of varying the number of households and the time frame on the optimal policies. It is observed that the policies remain consistent irrespective of these factors. Additionally, the stability and responsiveness of the optimal control policy are discussed concerning the MPC time horizon, with longer horizons providing more accurate predictions but requiring higher computational effort.

Future efforts can be directed towards incorporating on-the-fly uncertainties overlaid on top of a smoothly varying demand and solar power supply. Also, extracting a more regularized and implementable but efficient charging/discharging policy from the obtained optimal policy can be of great interest.

7 Video Presentation link and approval

The link to the video presentation: here

We approve of sharing this report with future students.

References

[1] Thomas Morstyn, Branislav Hredzak, Ricardo P Aguilera, and Vassilios G Agelidis. Model predictive control for distributed microgrid battery energy storage systems. *IEEE Transactions on Control Systems Technology*, 26(3):1107–1114, 2017.

[2] Qinglai Wei, Guang Shi, Ruizhuo Song, and Yu Liu. Adaptive dynamic programming-based optimal control scheme for energy storage systems with solar renewable energy. *IEEE Transactions on Industrial Electronics*, 64(7):5468–5478, 2017.

[3] Lukas Ortmann, Christian Rubin, Alessandro Scozzafava, Janick Lehmann, Saverio Bolognani, and Florian Dörfler. Deployment of an online feedback optimization controller for reactive power flow optimization in a distribution grid. *arXiv preprint arXiv:2305.06702*, 2023.

[4] Chenxiao Guan, Yanzhi Wang, Xue Lin, Shahin Nazarian, and Massoud Pedram. Reinforcement learning-based control of residential energy storage systems for electric bill minimization. In *2015 12th Annual IEEE Consumer Communications and Networking Conference (CCNC)*, pages 637–642. IEEE, 2015.

[5] Ying Li and Peter C Flynn. Deregulated power prices: comparison of diurnal patterns. *Energy policy*, 32(5):657–672, 2004.

[6] Akshay Agrawal, Robin Verschueren, Steven Diamond, and Stephen Boyd. A rewriting system for convex optimization problems. *Journal of Control and Decision*, 5(1):42–60, 2018.

[7] Steven Diamond and Stephen Boyd. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83):1–5, 2016.

[8] Nidhi Hegde, Laurent Massoulié, Theodoros Salonidis, et al. Optimal control of residential energy storage under price fluctuations. *Energy*, 2011.