The Potential for Building Emissions Reduction through Sporadic and Targeted Load Modulation

Utkarsha Agwan [1] Samuel Bobick [2] Srinath Rangan [2] Kameshwar Poolla [1] Costas Spanos [1]

Abstract

This paper focuses on pragmatic load flexibility measures for building electricity consumers that might not be able to modulate their load continuously, and instead act through short-term and sporadic load shift and shed, e.g., by shedding load a few times a year. Through a data-driven simulation of load modulation measures, we find that there are high-value time periods which account for a disproportionate amount of emissions, and consequently, even sporadic voluntary load shift and shed can achieve reductions in CO_2 emissions associated with electricity consumption. We quantify the emissions reductions for different load flexibility scenarios and discuss the practical implications associated with implementing sporadic and targeted load management measures in response to marginal emissions signals.

1. Introduction

The majority of greenhouse gas emissions caused by human activity are linked to the consumption of energy. The buildings sector is a significant part of this—energy consumption in residential and commercial buildings accounts for $\sim$17% of greenhouse gas emissions worldwide (Ritchie, 2020). Methods to reduce building emissions include electrifying end-uses such as heating, installing local clean generation such as solar, and using storage backup to increase self-consumption of onsite clean energy generation. Another way of reducing emissions is to reduce energy consumption overall, for example, through energy efficiency measures. Yet another method is to *modulate consumption specifically in the periods when the emissions impact is the highest*. This is the method that we will focus on.

Load modulation, i.e., *the practice of electricity consumers changing their energy consumption in response to grid signals*, is a valuable resource to accommodate the inherent variability of solar and wind generation and to preferentially use low emissions energy generation. Load modulation can entail load shift, where load is reduced at one time and increased at another; or load shed, where energy consumption is reduced during a time period without a subsequent load increase in some other time period. Any load modulation is a deviation from consumers' desired electricity consumption level and will result in some *disutility*. For loads which cannot be automated for reasons such as cost, infrastructure requirement, or privacy, implementing recurring load modulation can be difficult. However, if load modulation is not too frequent, it can be implemented manually. As evidenced in California's emergency call for demand response in Fall 2022 (St. John, 2022), there are many flexible loads which do not participate in recurring demand response markets, but are available for one-off load shifts and sheds.

Most work has centered around the machine learning, control, and optimization problems associated with utilizing grid signals to constantly modulate building electricity consumption at the household or device level (Bovornkeeratiroj et al., 2023; Lu et al., 2019; Wang et al., 2020; Wen et al., 2015). In addition to traditional building loads, prior studies have focused on loads such as electric vehicles (Bovornkeeratiroj et al., 2023; Daniels et al., 2022), batteries (Bovornkeeratiroj et al., 2023; Jha et al., 2020), and data centers (Bostandoost et al., 2024; Sukprasert et al., 2024). The authors of (Carmichael et al., 2021) provide a high-level evaluation of demand flexibility's role in emissions reduction, considering a wide array of decarbonization scenarios, grid operating conditions, and emissions signals used for control. To the best of our knowledge, no other work has analyzed the impact of *sporadic* load modulation on emissions. Sporadic manual load modulation offers a pragmatic control approach for loads that are not highly networked or "smart." Furthermore, the sporadic nature of our approach reduces disutility associated with more frequent or continuous interventions.

[1]Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA, USA [2]College of Computing, Data Science, and Society, University of California, Berkeley, CA, USA. Correspondence to: Samuel Bobick <samuelbobick@berkeley.edu>, Srinath Rangan <srinath.rangan@berkeley.edu>.

Proceedings of the 42nd International Conference on Machine Learning, Vancouver, Canada. PMLR 267, 2025.

This paper's contributions are as follows:

1. We propose load shift and shed optimization formulations that rely on the marginal emissions rate as a control signal.

2. Through a data-driven simulation, we calculate potential emissions reductions for a variety of customers and load flexibility scenarios using *targeted* (i.e., short duration) and *sporadic* (i.e., infrequent) load modulation.

3. We discuss the practical implications and future research directions for implementing targeted and sporadic load modulation.

2. Problem Formulation

2.1. Load Shift

Let $\boldsymbol{\lambda}$ denote the time series of marginal emissions intensities (in lbs. CO_2/kWh), and let $\mathbf{d}$ represent the consumer's load curve (in kWh). The consumer minimizes total emissions by choosing $\mathbf{d}$:

$$\min_{\mathbf{d}} \sum_{t=1}^{T} \lambda(t)d(t). \tag{1}$$

The consumer's uncontrolled (i.e., utility-maximizing) load profile is denoted by $\hat{\mathbf{d}}$. To protect utility, we constrain how much $\mathbf{d}$ may deviate from $\hat{\mathbf{d}}$, based on three flexibility parameters:

- γ: fraction of load at timestep $\hat{t}$ that may be shifted;

- Δ: length, in hours, of the window over which shifted load may be redistributed;

- ν: kick-back ratio which determines the additional energy that must be consumed because due to the shift.

We assume the consumer curtails load at a single timestep $\hat{t}$. This load is redistributed to other timesteps within a window of size 2Δ centered at $\hat{t}$. In our simulation, $\hat{t}$ is chosen such that the window $[\hat{t}-\Delta, \hat{t}+\Delta]$ has the widest range of marginal emissions possible.

The curtailment occurs only at $\hat{t}$, where energy consumption is reduced by a factor of $(1-\gamma)$. For some loads (e.g. heating), this curtailed energy must be compensated for by increased consumption in neighboring timesteps. This redistributed energy is scaled by the kick-back factor ν.

The resulting constraints are:

$$d(t) \geq \hat{d}(t), \quad \forall \quad t \neq \hat{t}, \tag{2}$$

$$d(\hat{t}) = (1-\gamma)\hat{d}(\hat{t}), \tag{3}$$

$$\sum_{[t=\hat{t}-\Delta, \hat{t}+\Delta] \setminus \{\hat{t}\}} \frac{d(t)-\hat{d}(t)}{1+\nu|t-\hat{t}|} = \gamma \hat{d}(\hat{t}). \tag{4}$$

2.2. Load Shed

We model a consumer capable of shedding a proportion γ of their load over a time window of length Δ. Let $\hat{t}$ denote the start of the load shed window.

We select $\hat{t}$ by solving:

$$\hat{t} = \arg\max_{i} \sum_{t=i}^{i+\Delta} \lambda(t). \tag{5}$$

The consumer then reduces load within the interval $[\hat{t}, \hat{t}+\Delta]$ by a factor of $(1-\gamma)$:

$$d(t) = \begin{cases} (1-\gamma)\hat{d}(t), & \text{if } t \in [\hat{t}, \hat{t}+\Delta] \\ \hat{d}(t), & \text{otherwise.} \end{cases} \tag{6}$$

3. Numerical Study

3.1. Data

3.1.1. Energy Consumption

We use NREL's ComStock and ResStock datasets (Frick et al., 2019), which model commercial and residential energy use disaggregated by end use based on detailed building simulations. Specifically, we use the county-level aggregated time series for medium and large commercial customers in Alameda County, CA simulated using meteorological data from 2018, available at 15-minute resolution.

3.1.2. Marginal Emissions Intensity

We use a dataset of estimated marginal emissions intensity from WattTime (WattTime, 2022) to obtain historical marginal operating emissions rates (MOERs) for the Northern California grid. This data is available at 5-minute intervals, but we downsample to 15-minute intervals to match the granularity of the energy consumption dataset.

3.2. Simulation Results

We simulate the load shift and shed formulations from Section 2 for the four largest load categories in ComStock and

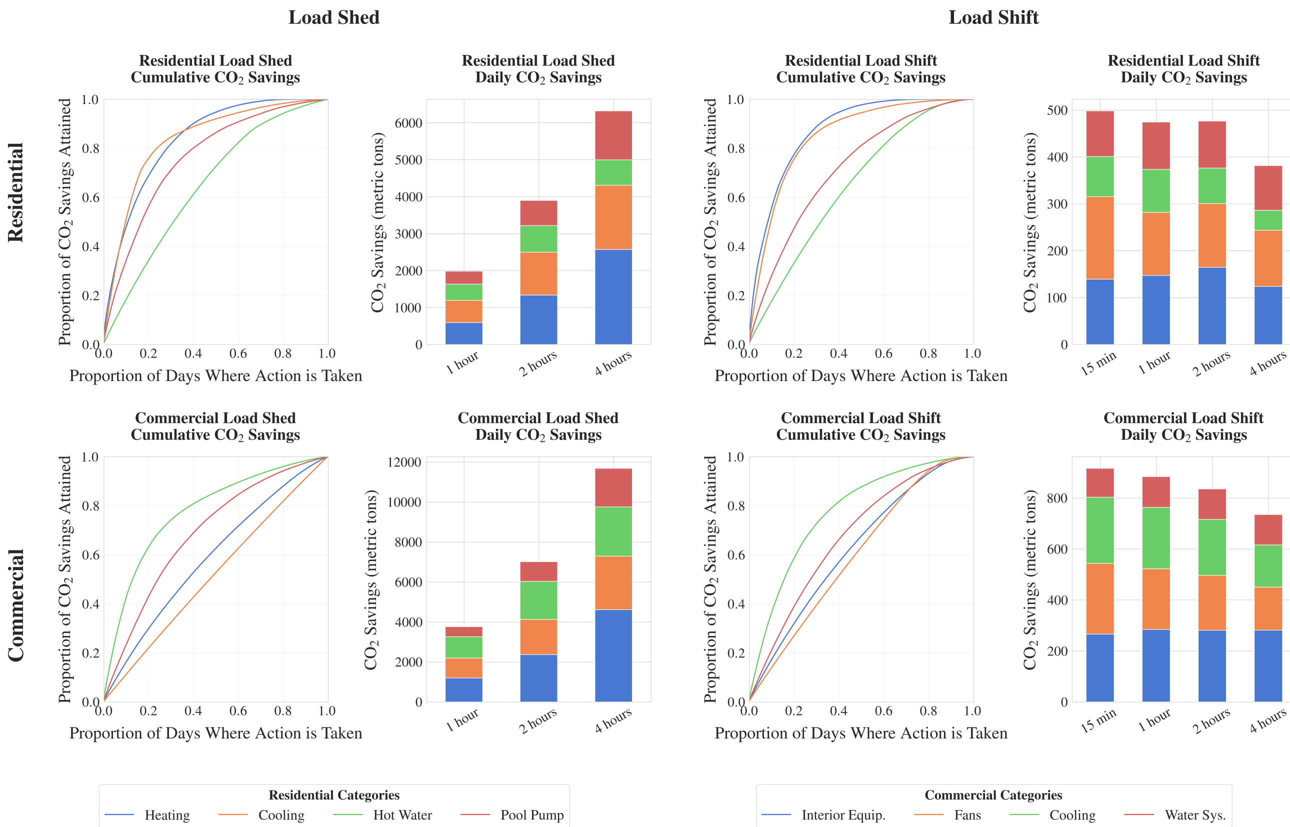

Figure 1. The left panel of each subplot shows the cumulative effect of load modulation for a 2-hour flexibility assumption when taking actions in the order of magnitude of emissions reduction. The right panel shows the amount of CO_2 savings from taking the 30 actions with the highest propensity to reduce emissions across different assumptions of Δ.

ResStock.[1] We focus our analysis on Alameda County in California, USA. Additional simulation details are described in Appendix A.

Since each intervention causes some disutility to electricity consumers, loads with a good propensity for load modulation will have a Pareto distribution, where taking action for a small proportion of the days of the year will capture most of the possible CO_2 savings. Figure 1 shows this property. This property is most prominent in residential heating and cooling—nearly 80% of emissions savings can be achieved by taking a single load modulation action on 20% of the days of the year.

4. Discussion

Our work explores a specific hypothesis: consumers can achieve significant emissions reductions even with sporadic load modulation. This widens the scope of impact for marginal emissions signals by inducing additional customer categories to engage in load modulation, as opposed to the EV, battery, and data center customer categories which have been the primary user of emissions intensity data. By specifically targeting sporadic modulation, we offer a pragmatic pathway to decrease disutility and unlock emissions reduction potential in building loads that may not be amenable to such continuous interventions. There are several practical considerations to keep in mind while implementing load modulation mechanisms.

The Effect of Forecast Quality on Emissions Savings

In order to maximize the impact of a few, infrequent interventions, it is important to have robust forecasting and signaling technology to implement load modulation. Accurate and timely forecasts are essential, particularly if marginal emissions rates vary rapidly over the course of a day, since a mistimed shift action could actually end up increasing emissions. Similarly, electricity customers must have access to the live marginal emissions status of the power grid, so that they can optimize their energy consumption at finer time scales. While some market operators such as (PJM, 2022)

[1] Code available at https://github.com/utkarshapets/load-shift-emissions.

publish marginal emissions data, most of these analyses are post-hoc. Non-market entities such as (WattTime, 2022) and (Electricity Maps, 2025) can fill this role in the future. However, the relationship between marginal emissions forecast quality and load modulation efficacy remains an open question.

Location Matters

The emissions intensity is determined by the energy mix of the grid, and each power grid has a different generation source composition. Even within a power network, line constraints may impact the marginal generator, introducing the idea of a *locational* marginal emissions intensity. This is a potential future extension of our work.

The Energy Mix of the Future Looks Different

The grid of the future will have a larger share of renewables, and the power grid in California could be a good proxy for this. The authors of (Agwan et al., 2023) evaluate the marginal emissions rate as a signal for load management and find that WattTime MOERs in Northern California often exhibit substantial intra-day variation. This suggests that shifting load within a day has a high potential for reducing emissions.

Machine Learning for Sporadic Load Modulation

Methodologies from machine learning will be useful for developing targeted load modulation mechanisms. For example, marginal emissions forecasts used to identify optimal intervention periods such as (WattTime, 2022) are grounded in machine learning. AI-driven optimization approaches such as reinforcement learning, hold potential for developing adaptive and personalized sporadic load shift strategies that learn optimal intervention periods and actions, effectively responding to grid conditions (e.g. emissions and price signals) and user preferences.

References

Agwan, U., Bobick, S., Rangan, S., Poolla, K., and Spanos, C. J. Time varying marginal emissions intensity of energy consumption: Implications for flexible loads. *Findings*, 2023.

Alstone, P., Potter, J., Piette, M. A., Schwartz, P., Berger, M. A., Dunn, L. N., Smith, S. J., Sohn, M. D., Aghajanzadeh, A., Stensson, S., et al. 2025 california demand response potential study-charting california's demand response future. final report on phase 2 results. Technical report, Lawrence Berkeley National Lab.(LBNL), Berkeley, CA (United States), 2017.

Bostandoost, R., Lechowicz, A., Hanafy, W. A., Bashir, N., Shenoy, P., and Hajiesmaili, M. Lacs: Learning-augmented algorithms for carbon-aware resource scaling with uncertain demand. In *Proceedings of the 15th ACM International Conference on Future and Sustainable Energy Systems*, pp. 27–45, 2024.

Bovornkeeratiroj, P., Bashir, N., Deulkar, V., Balaji, B., Shenoy, P., Irwin, D., and Hajiesmaili, M. Quantifying the decarbonization potential of flexible load. In *Proceedings of the 10th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*, pp. 429–433, 2023.

Carmichael, C., Mandel, J., Richardson, H., Taylor, E., and Usry, C. The carbon emissions impact of demand flexibility. *RMI. https://rmi. org/our-work/buildings*, 2021.

Daniels, L., Gross, B., Lewis, C., and Stone, L. More evs, fewer emissions: How to maximize emissions reductions by smart charging electric vehicles. *RMI. http://www.rmi.org/insight/more-evs-fewer-emissions*, 2022.

Electricity Maps. Electricity maps api documentation. Technical report, Electricity Maps, 2025. URL https://portal.electricitymaps.com/docs.

Frick, N. M., Wilson, E., Reyna, J., Parker, A., Present, E., Kim, J., Hong, T., Li, H., and Eckman, T. End-use load profiles for the us building stock: Market needs, use cases, and data gaps. 2019.

Jha, R., Lee, S., Iyengar, S., Hajiesmaili, M. H., Irwin, D., and Shenoy, P. Emission-aware energy storage scheduling for a greener grid. In *Proceedings of the Eleventh ACM International Conference on Future Energy Systems*, pp. 363–373, 2020.

Lu, R., Hong, S. H., and Yu, M. Demand response for home energy management using reinforcement learning and artificial neural network. *IEEE Transactions on Smart Grid*, 10(6):6629–6639, 2019.

PJM. Marginal emissions rates - a primer. Technical report, PJM, 2022. URL https://www.pjm.com/-/media/etools/data-miner-2/marginal-emissions-primer.ashx.

Ritchie, H. Sector by sector: where do global greenhouse gas emissions come from? *Our World in Data*, 2020. https://ourworldindata.org/ghg-emissions-by-sector.

St. John, J. Californians saved the grid again. they should be paid more for it. 2022. URL https://www.canarymedia.com/articles/grid-edge/californians-saved-the-grid-again-%they-should-be-paid-more-for-it.

Sukprasert, T., Souza, A., Bashir, N., Irwin, D., and Shenoy, P. On the limitations of carbon-aware temporal and spatial workload shifting in the cloud. In *Proceedings of the Nineteenth European Conference on Computer Systems*, EuroSys '24, pp. 924–941, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400704376. doi: 10.1145/3627703.3650079. URL https://doi.org/10.1145/3627703.3650079.

Wang, B., Li, Y., Ming, W., and Wang, S. Deep reinforcement learning method for demand response management of interruptible load. *IEEE Transactions on Smart Grid*, 11(4):3146–3155, 2020.

WattTime. Marginal emissions modeling: WattTime's approach to modeling and validation. Technical report, WattTime, 2022. URL https://www.watttime.org/app/uploads/2022/10/WattTime-MOER-modeling-20221004.pdf.

Wen, Z., O'Neill, D., and Maei, H. Optimal demand response using device-based reinforcement learning. *IEEE Transactions on Smart Grid*, 6(5):2312–2324, 2015.

A. Simulation Details

We base our simulation off of ComStock and ResStock's load curves derived from the 2018 meterological year, along with Wattime's 2018 MOER values for the CAISO-North region. Table 1 shows the flexibility assumptions used in the simulation. We base our flexibility assumptions on those presented in (Alstone et al., 2017), with slight modifications. In particular, we use (Alstone et al., 2017)'s load flexibility framework that models direct load control (i.e., manual intervention rather than continuous control). Categories represent individual load curves modeled in the ComStock and ResStock datasets. Furthermore, borrowing from (Alstone et al., 2017), we penalize heating cooling, and water heating with a 5% energy penalty per hour of shift. That is, we set $\nu = 0.05$. For all other loads, we set $\nu = 0$.

Building Type	**Category**	$\gamma \mid \Delta$ **= 0.25**	$\gamma \mid \Delta$ **= 1**	$\gamma \mid \Delta$ **= 2**	$\gamma \mid \Delta$ **= 4**
Commercial	Interior Equipment	0.35	0.35	0.35	0.35
Commercial	Fans	0.60	0.50	0.45	0.35
Commercial	Cooling	0.60	0.50	0.45	0.35
Commercial	Water Systems	0.84	0.84	0.84	0.84
Residential	Heating	0.60	0.40	0.40	0.35
Residential	Cooling	0.60	0.40	0.40	0.35
Residential	Hot Water	0.95	0.95	0.80	0.40
Residential	Pool Pump	0.79	0.70	0.70	0.70

Table 1. γ, the proportion of load that can be modulated, for different values of Δ, the size of the modulation window, for selected commercial and residential load categories.