EDGIE: A simulation testbed for investigating the impacts of building and vehicle electrification on distribution grids

Priyadarshan
School of Mechanical Engineering,
Purdue University, USA
priyada@purdue.edu

Elias N. Pergantis
School of Mechanical Engineering,
Purdue University, USA
epergant@purdue.edu

Constance Crozier
School of Industrial
and Systems Engineering,
Georgia Institute of Technology, USA
ccrozier8@gatech.edu

Kyri Baker
Dept. of Civil, Architectural
and Environmental Engineering,
University of Colorado Boulder, USA
kyri.baker@colorado.edu

Kevin J. Kircher
School of Mechanical Engineering,
Purdue University, USA
kircher@purdue.edu

Abstract

Replacing fossil-fueled appliances and vehicles with electric versions can significantly reduce emissions. However, electric heating and vehicle charging can cause peaks in electricity demand that stress infrastructure in buildings and power grids, jeopardizing reliability or forcing costly infrastructure upgrades. This paper presents the open-source EDGIE (Emulating the Distribution Grid Impacts of Electrification) toolbox. EDGIE matches experiment data from an all-electric home in cold weather. It can simulate many locations and levels of technology adoption, and supports optimization and network power flow simulation. In simulations of a fully electrified neighborhood during the coldest week of 2019 in New York City, demand peaks at quadruple today's summer peak. Peaks are particularly sensitive to the use of overnight thermostat set-point reductions and to the efficiencies of heat pumps and building envelopes. Optimal vehicle-to-home coordination with flexible space and water heating reduces peak demand by 35% and transformer degradation by 99%.

Keywords: Electrification, Distribution Grid, Heat Pump, Electric Vehicle, Transformer.

1. Introduction

Burning fossil fuels in power plants, residential and commercial buildings, and light-duty passenger vehicles causes 54% of United States greenhouse gas emissions (EPA (2021)). Decarbonizing electricity generation and replacing fossil-fueled vehicles and appliances with electric versions could essentially eliminate these emissions. However, widespread electrification could cause high peaks in electricity demand, particularly on the coldest days. These peaks could stress electrical infrastructure in low-voltage building networks, medium-voltage distribution networks, or high-voltage transmission networks. Without infrastructure upgrades or peak demand mitigation strategies, electrification could cause voltages to drop below acceptable levels. It could also overheat power lines or transformers, risking equipment failures or blackouts. On the other hand, overly conservative infrastructure upgrades could cost ratepayers billions of dollars in unnecessary price increases, disproportionately burdening people in poverty. Better understanding how electrification might impact the grid, and how to mitigate those impacts, could reduce these risks to reliability and affordability.

Several recent studies have explored the potential impacts of electrification on United States power systems. Blonsky et al. (2019) estimate that electric vehicles may use 551 TWh of electricity by 2040, compared to 1 TWh in 2017. Scenarios modeled by EPRI (2018) show electrification increasing annual electricity use by 24–52% over 2015 levels in 2050. Zhou and Mai (2021) use PLEXOS, commercial software for power system optimization, to assess the value of demand flexibility under scenarios with varying degrees of electrification. They find that demand flexibility could reduce annual electricity production costs in 2050 by about 10%. Tarroja et al. (2018) use HIGRID, a simulation toolbox based on EnergyPlus prototype models, to assess the impacts of heating electrification in California. They estimate a 32% peak demand increase in 2050 relative to 2015. White et al. (2021) use EnergyPlus to retrospectively model space heating electrification in Texas in 2016. They find that peak demand would have increased by 10 GW, about 14% of the Texas grid's 2016 peak. Elmallah et al. (2022) model the impacts of home and vehicle

URI: https://hdl.handle.net/10125/106760
978-0-9981331-7-1

electrification on the Pacific Gas & Electric distribution network in northern California. They estimate $1–10 billion in required distribution grid upgrades between now and 2050 across a range of electrification scenarios.

While the above studies demonstrate some potential impacts of electrification on the grid, important questions remain unanswered. For example, the above studies focus mainly on regions with relatively mild winters, such as California and Texas. Grid impacts in cold climates may differ substantially. With the exception of Elmallah et al. (2022), the above studies all focus on the transmission level. Distribution-level considerations, such as voltage regulation and transformer degradation, have received less attention. Other than Zhou and Mai (2021), the above studies focus mainly on characterizing the problem, rather than developing solutions. To develop solutions, it helps to have a simulation testbed that can quickly run sensitivity analyses to reveal the key drivers of grid impacts, and that interfaces easily with optimization software.

This paper introduces an open-source simulation testbed (https://github.com/priyada7/EDGIE) aimed at helping researchers investigate these questions. The testbed, named EDGIE (Emulating the Distribution Grid Impacts of Electrification), can simulate distribution networks of varying sizes and topologies in a wide range of locations and technology adoption scenarios. It models the thermal dynamics and degradation of transformers, key distribution grid components. EDGIE models buildings, space and water heating systems, and electric vehicle batteries using linear differential equations that capture their essential dynamics. Solving these equations analytically reduces simulation to matrix-vector multiplication, which basic linear algebra routines can execute very efficiently. This efficiency enables rapid parameter sweeps and other sensitivity analyses that can reveal key drivers of grid impacts. EDGIE also interfaces with packages for convex optimization and power flow simulation.

Understanding the impacts of winter peak loads on distribution infrastructure requires accurate models of component dynamics. EDGIE incorporates models which include outdoor temperature dependency for heat pump capacities and coefficients of performance (COPs), for electric vehicle driving efficiencies, and for transformer degradation. This means that for modeling distribution systems under extreme temperatures, EDGIE can provide a more accurate picture of component operation compared to existing testbeds. EDGIE's space heating models closely match real data from an all-electric home in very cold weather.

This paper presents the models and input data that EDGIE uses. It demonstrates EDGIE's simulation, sensitivity analysis, and optimization capabilities through the example of about 1,000 homes in New York on a very cold week. Under 100% electrification with air-to-air heat pumps and resistance backup heating, the neighborhood's winter demand peaks at 16.6 MW, quadruple today's summer peak. The winter peak depends sensitively on the heat pump sizing approach, on whether people reduce their space heating temperature set-points overnight, on home sizes and insulation values, and on heat pumps' cold-weather COPs. In an optimization example, coordinating electric vehicle charging and discharging with flexible space and water heating reduces the winter peak by 5.8 MW (35%) and transformer degradation by 99%.

2. Models and Input Data

2.1. Base Electricity Demand

EDGIE includes electricity demand data for each home in the base case of fossil-fueled driving and heating. The electricity demand data come from the Multi-Family Residential Electricity Dataset, compiled by Meinrenken et al. (2020). This dataset includes time-series measurements from 390 apartments in New York City in 2019, averaged over groups of 15 apartments to protect user privacy. The measurements include electricity used for nearly all purposes, including cooling. However, they do not include electricity for space or water heating, as the apartment buildings burned natural gas for heat. To generate each home's base electricity load, EDGIE randomly samples a load profile from one group of apartments, then re-scales the data according to floor area. In simulations of electrified neighborhoods, EDGIE adds the power used for electric vehicle charging, water heating and space heating to the base electricity demand. EDGIE uses low-order linear differential equation models based on Kircher et al. (2021). These models appear widely in the research literature.

2.2. Electric Vehicles

EDGIE models electric vehicle batteries via

$$\begin{aligned} &\dot{E}(t) = -rE(t) + \eta_1 p_1(t) - w_1(t) \\ &0 \le E(t) \le \overline{E} \\ &0 \le p_1(t) \le \overline{p}_1. \end{aligned} \tag{1}$$

Here t (h) denotes time, E (kWh) is the chemical energy stored in the battery, r (1/h) is the battery's self-dissipation rate, η_1 is the charging efficiency, p_1

(kW) is the electric charging power, w_1 (kW) is the chemical power discharged to drive the vehicle, $\overline{E}$ (kWh) is the energy capacity, and $\overline{p}_1$ (kW) is the charging power capacity. Assuming a zero-order hold on the input signals over each time step of duration Δt (h) gives the exact discrete-time dynamics,

$$E(k+1) = a_1 E(k) + \frac{1-a_1}{r}(\eta_1 p_1(k) - w_1(k)), \quad (2)$$

where k indexes time steps and $a_1 = \exp(-r\Delta t)$.

In the example simulations in this paper, each vehicle uses a battery with capacity $\overline{E}$ = 50–82 kWh, consistent with manufacturer data on a range of popular vehicles. (EDGIE randomly generates all input data mentioned with ranges, such as 50–82 kWh, from independent symmetric triangular distributions.) The vehicles charge on either Level 1, modeled as a 120 V/15 A circuit, or Level 2 (240 V/48 A), giving charge capacities $\overline{p}_1$ of 1.8 or 11.5 kW. Round-trip commutes range from 15–35 km per day, consistent with DOT (2017) data. The vehicles have charging efficiencies η_1 of 0.9–0.95, consistent with manufacturer data (Novac et al. (2013)). EDGIE computes the discharged power w_1 by multiplying the commute distance times the vehicle's energy intensity (in kWh/km).

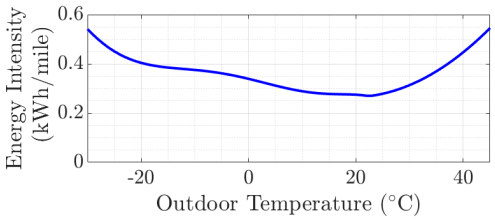

Figure 1. The simulated vehicles use more energy per unit distance in very cold or very hot weather.

An electric vehicle's energy intensity (the amount of energy it discharges to drive a given distance) depends on the outdoor air temperature. It increases in both cold and hot weather. EDGIE captures these effects through a model developed by Yuksel and Michalek (2015) and illustrated in Fig. 1. In simulations, each user recharges their battery when its state of charge drops below 15–25%, consistent with some manufacturer recommendations. To recharge, users plug their vehicles in when they arrive home at 6–10 PM. The vehicles charge at maximum power until their batteries are full.

2.3. Water Heating

EDGIE models water heaters via

$$\begin{aligned} C_2\dot{T}_2(t) &= \frac{\theta_2 - T_2(t)}{R_2} + q_2(t) - w_2(t) \\ 0 &\le q_2(t) \le \eta_2\overline{p}_{2h} + \overline{p}_{2r}. \end{aligned} \quad (3)$$

Here T_2 (°C) is the water temperature, C_2 (kWh/°C) is the water's thermal capacitance, θ_2 (°C) is the (constant) air temperature surrounding the tank, R_2 (°C/kW) is the thermal resistance between the water and surrounding air, q_2 (kW) is the thermal power supplied to the tank, w_2 (kW) is the thermal power withdrawn for showers, dish-washing, laundry, *etc.*, η_2 is the heat pump's COP, $\overline{p}_{2h}$ (kW) is the heat pump's electric power capacity, and $\overline{p}_{2r}$ (kW) is the electric power capacity of the resistance heater. In discrete time,

$$\begin{aligned} T_2(k+1) = a_2 T_2(k) + (1-a_2)[\theta_2 \\ + R_2(q_2(k) - w_2(k))], \end{aligned} \quad (4)$$

where $a_2 = \exp(-\Delta t/(R_2C_2))$.

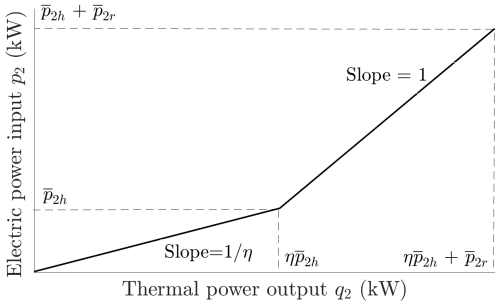

Figure 2. The control logic prioritizes the more efficient heat pump, adding resistance if needed.

This model can represent a water heater with a heat pump only (if $\overline{p}_{2r} = 0$), electric resistance only (if $\overline{p}_{2h} = 0$), or a hybrid of the two (if $\overline{p}_{2h} > 0$ and $\overline{p}_{2r} > 0$). Fig. 2 illustrates the hybrid case, wherein the control logic prioritizes the more efficient heat pump and supplements it with resistance heating only if the heat pump cannot meet thermal load alone. In general, supplying thermal power $q_2(t)$ to the water uses electric power $p_2(t)$, given by the following piecewise definition:

$$\begin{cases} 0 & q_2(t) \le 0 \\ q_2(t)/\eta_2 & 0 < q_2(t) \le \eta_2\overline{p}_{2h} \\ (1-\eta_2)\overline{p}_{2h} + q_2(t) & \eta_2\overline{p}_{2h} < q_2(t) \le \overline{p}_{2h} + \overline{p}_{2r} \\ \overline{p}_{2h} + \overline{p}_{2r} & \overline{p}_{2h} + \overline{p}_{2r} < q_2(t). \end{cases} \quad (5)$$

In simulations, each water heater has a user-specified temperature set-point $\hat{T}_2(k)$ at each time k. Driving the water temperature from the current state $T_2(k)$ to the next set-point $\hat{T}_2(k+1)$ would require thermal power

$$\hat{q}_2(k) = \frac{1}{R_2}\left(\frac{\hat{T}_2(k+1) - a_2 T_2(k)}{1-a_2} - \theta_2\right) + w_2(k). \tag{6}$$

Each water heater's control system attempts to perfectly track its set-point by delivering the thermal power $\hat{q}_2(k)$, but may saturate at an upper or lower capacity limit:

$$q_2(k) = \max(0, \min(\eta_2 \overline{p}_{2h} + \overline{p}_{2r}, \hat{q}_2(k))). \tag{7}$$

The cylindrical water heaters simulated here have 50–80 gallon (0.23–0.36 m^3) volumes and 10 inch (0.25 m) radii, typical dimensions for single-family homes. Each thermal capacitance is the product of the density, specific heat and volume of the water in the tank. Each thermal resistance follows from dividing an R-value of 6–8 °F ft^2/BTU/h (1060-1410 °C m^2/kW) by the tank's vertical surface area. Water withdrawals come from the Building America Domestic Hot Water Event Schedule Generator, which generates plausible hot water withdrawals using the method in Hendron et al. (2010). The temperature of inlet water to the tank varies sinusoidally, with mean 10 °C, amplitude 5 °C, peak on August 5, and period of one year. This inlet water temperature profile follows the empirical data reported in Burch and Thornton (2012). The other input parameters can be found in Kircher et al. (2021).

2.4. Space Heating

This study models space heating using similar methods to those used for water heating. Each home has an air-to-air heat pump, possibly supplemented with backup resistance heating. A first-order linear model captures the home's temperature dynamics:

$$\begin{aligned} C_3 \dot{T}_3(t) &= \frac{\theta_3(t) - T_3(t)}{R_3} + q_3(t) + w_3(t) \\ 0 &\le q_3(t) \le \eta_3(\theta_3(t))\overline{p}_{3h} + \overline{p}_{3r}. \end{aligned} \tag{8}$$

Here T_3 (°C) is the indoor air temperature, θ_3 (°C) is the outdoor air temperature, C_3 (kWh/°C) is the thermal capacitance of the indoor air and any tightly-coupled thermal mass, R_3 (°C/kW) is the thermal resistance between the indoor and outdoor air, q_3 (kW) is the thermal power supplied by the heat pump and/or resistance heater, w_3 is the exogenous thermal power from the sun, plug loads, lights, body heat, *etc*., η_3 is the heat pump's COP, $\overline{p}_{3h}$ (kW) is the heat pump's electric power capacity, and $\overline{p}_{3r}$ (kW) is the electric resistance heating element's power capacity. In discrete time,

$$\begin{aligned} T_3(k+1) = a_3 T_3(k) + (1-a_3)[\theta_3(k) \\ + R_3(q_3(k) + w_3(k))], \end{aligned} \tag{9}$$

where $a_3 = \exp(-\Delta t/(R_3 C_3))$.

In simulations, each heat pump has a user-specified temperature set-point $\hat{T}_3$ that may vary with time. Each heat pump's control system either tracks its set-point exactly or saturates at a capacity limit. Driving the indoor air temperature from $T_3(k)$ to the next set-point $\hat{T}_3(k+1)$ would require thermal power similar to Eq. (6) with appropriate subscript changes. As with water heating, the space heating control logic prioritizes the heat pump and uses resistance heating only if the heat pump cannot meet thermal load alone:

$$q_3(k) = \max(0, \min(\eta_3(\theta_3(k))\overline{p}_{3h} + \overline{p}_{3r}, \hat{q}_3(k))). \tag{10}$$

Providing the thermal power $q_3(k)$ uses total electric power $p_3(k)$, given by a piecewise definition that matches the water heater definition (5) up to the appropriate subscript changes.

The simulations in this paper use heat pump and resistance heater capacities of $\overline{p}_{3h} = 4.5$ kW and $\overline{p}_{3r} = 18$ kW. The heat pump COPs vary linearly with the outdoor temperature, with slopes and intercepts tuned to specification sheets from several manufacturers. COPs are 1.5–2 at -15 °C and 3–4 at 7 °C. Homes have 200–250 m^2 of total floor area. EDGIE calculates the thermal resistance R_3 by dividing an overall R-value of 2.8–3.9 °F ft^2/BTU/h (500–700 °C m^2/kW) by the home's total exterior wall area, assuming a square footprint and two stories of 3–3.6 m height. The overall R-value accounts for insulation, framing, windows, doors, infiltration of outdoor air, *etc.* This process results in a mean R_3 of about 2.1 °C/kW. The thermal capacitance C_3 follows from the volume, density, and specific heat capacity of the indoor air; it varies from 1-3 (kWh/°C). EDGIE computes w_3 by rescaling the global solar irradiance on a horizontal surface (which most weather data services provide), adding the base electricity demand (to model heat from lights, electronics, *etc.*), and adding random noise representing body heat. The resulting exogenous thermal power intensity ranges from about 8–12 W per m^2 of floor area.

2.5. Space Heating Experimental Validation

Simulations later in this paper suggest that space heating is the primary driver of demand peaks from electrification in cold climates. This makes validating

EDGIE's space heating model particularly important. This paper therefore validates the space heating model experimentally against an all-electric home in a cold climate (West Lafayette, Indiana, USA). The validation data span December 10–31, 2022, when a severe cold snap brought outdoor air temperatures below -20 °C. This cold snap allows model validation under nearly worst-case conditions.

Figure 3. The experimental testbed, a 1920s-era home retrofit with all electric appliances.

Fig. 3 shows the experimental testbed, an all-electric home near Purdue University's campus. The 1920s-era home has two stories and 208 m^2 of floor area. Its exterior walls have R-21 °F ft^2/BTU/h (3,500 °C m^2/kW) foam insulation. Accounting for framing, as well as the U-1.4 BTU/h/ft^2/°F (8 W/m^2/°C) windows that make up about 20% of the exterior wall area, gives an overall R-value of about 3 °F ft^2/BTU/h (530 °C m^2/kW). The home has an air-to-air heat pump with four tons (14 kW) of rated cooling capacity and a Seasonal Energy Efficiency Ratio of 18 BTU/Wh (5.3 seasonal COP). The home has 19.2 kW of resistance heat.

Table 1 shows the validation results. The measured and simulated total electrical energy used for heating agree within 5% error. EDGIE overestimates the heat pump energy and underestimates the resistance energy, each by about 10%. EDGIE accurately matches the individual peak power measurements for both the heat pump and resistance, but overestimates the combined peak by about 5%. This overestimation makes the peak demand simulation results in this paper slightly conservative.

Table 1. Experiment vs. Simulation Results

	Heat Pump		Resistance		Combined	
	Exp't	Sim.	Exp't	Sim.	Exp't	Sim.
Peak kW	4.5	4.5	14.4	14.6	18.2	19.1
Total kWh	1,053	1,168	447	402	1,500	1,570

2.6. Transformers

Load changes due to electrification may stress transformers, conductors, capacitor banks, switch-gear, or other distribution grid components. Of these, transformers are among the most sensitive to demand peaks and the most expensive to replace (NREL (2023)). The literature contains several standards for modeling transformer degradation. All focus on the mean temperature of the top layer of oil inside the transformer, θ_{to} (K), and the temperature of the hottest spot within that layer, θ_{hs} (K). This paper uses the IEC 60076-7 standard from Swift et al. (2001), which models the temperature dynamics as

$$\begin{aligned}\tau_{\text{to}}\dot{\theta}_{\text{to}}(t) &= \theta_{\text{a}}(t) - \theta_{\text{to}}(t) \\ &\quad + (\hat{\theta}_{\text{to}} - \hat{\theta}_{\text{a}})\left(\frac{1+\rho I_{\text{pu}}(t)^2}{1+\rho}\right)^n \\ \tau_{\text{hs}}(\dot{\theta}_{\text{hs}}(t) - \dot{\theta}_{\text{to}}(t)) &= \theta_{\text{to}}(t) - \theta_{\text{hs}}(t) \\ &\quad + (\hat{\theta}_{\text{hs}} - \hat{\theta}_{\text{to}})I_{\text{pu}}(t)^{2m}.\end{aligned} \tag{11}$$

Here τ_{to} (h) is the top oil time constant, θ_{a} (K) is the ambient air temperature, ρ is the transformer's loss-of-load ratio, I_{pu} is the per-unit current, n is the top oil exponent, τ_{hs} (h) is the hot spot time constant, and m is the hot spot exponent. Variables with hats, such as $\hat{\theta}_{\text{a}}$, denote values at the transformer's rated conditions. The simulations in this paper use the parameter values in Table 2, which come from Elmoudi et al. (2006).

Table 2. Transformer Parameters

Parameter	Value
Rated power	16 MVA
Cooling modes	ONAF
Top oil time constant, τ_{to}	2.67 h
Rated top oil temperature rise, $\hat{\theta}_{\text{to}} - \hat{\theta}_{\text{a}}$	50.6 K
Loss-of-load ratio, ρ	9.09
Top oil exponent, n	0.9
Hot spot time constant, τ_{hs}	0.1 h
Rated hot spot temperature rise, $\hat{\theta}_{\text{hs}} - \hat{\theta}_{\text{to}}$	26 K
Hot spot exponent, m	0.8

Given the hot spot temperature, Swift et al. (2001) model the transformer's loss of life between times t_1 and t_2 (h) as

$$\int_{t_1}^{t_2} \exp\left[(15000\text{K})\left(\frac{1}{383\text{K}} - \frac{1}{\theta_{\text{hs}}(t)}\right)\right] \text{d}t. \tag{12}$$

If the hot spot temperature remains constant at $\theta_{\text{hs}} = 383$ K between t_1 and t_2, the transformer loses $t_2 - t_1$ hours of life. Higher values of θ_{hs} cause exponentially faster transformer degradation.

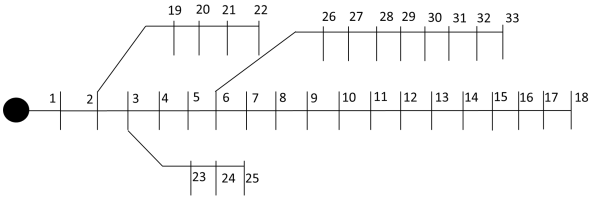

Figure 4. The IEEE 33 bus network topology.

2.7. Distribution Networks

To evaluate the impacts of electrification on distribution network voltages, EDGIE interfaces with MATPOWER (Zimmerman and Murillo-Sanchez (2020)), an open-source toolbox for power flow simulation and optimization. MATPOWER works with a wide variety of transmission and distribution networks. The simulations in this paper use the IEEE 33 bus distribution network shown in Fig. 4, with each bus corresponding to 30 homes. EDGIE computes the power draws at each bus in each time step, then feeds these values to MATPOWER, which solves the network's governing equations for the bus voltages.

3. Simulation Results for the Central Case

This section presents simulation results for 990 homes (30 per bus in the IEEE 33 bus network) during the coldest week of 2019 in New York City, when outdoor air temperatures dropped below -16 °C. This section discusses the central case, wherein all homes have air-to-air heat pumps and resistance backup for space heating, resistance water heaters, and electric vehicles with Level 2 chargers. The heat pumps, sized for cooling, run concurrently with resistance backup in the coldest hours. Users keep their indoor air temperature set-points constant, rather than reducing them overnight. Each home has two electric vehicles, consistent with the United States average of 1.86 (DOT (2017)). Simulating 990 homes over five days at one-hour time steps takes 12 seconds on a 2.9 GHz processor with 16 GB of RAM. The next section of this paper presents sensitivity analyses that perturb various parameters away from the central case described above.

Fig. 5 shows the load composition in the central case. The blue curve represents a 'no electrification' baseline, where all homes have fossil-fueled vehicles and heaters. Adding resistance water heaters (magenta curve) increases winter peak demand modestly, from about 3.4 MW to just above today's 3.9 MW summer peak (dashed red line). Adding electric vehicles (green curve), however, raises the peak over 7 MW, about 85% higher than today's summer peak. Adding heat pumps

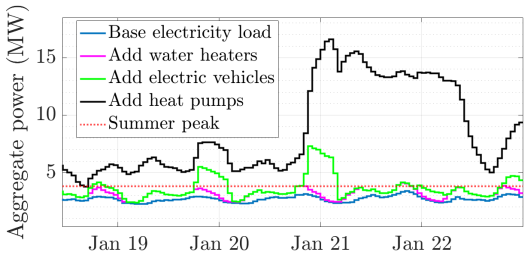

Figure 5. Aggregate load on the coldest days of 2019 in New York. Under full electrification, winter demand peaks at about 4 times today's summer peak.

with resistance backup (black curve) raises the peak to 16.5 MW on January 21, when the outdoor temperature dropped to -16 °C. This is roughly quadruple today's summer peak. Fig. 6 further illustrates these trends.

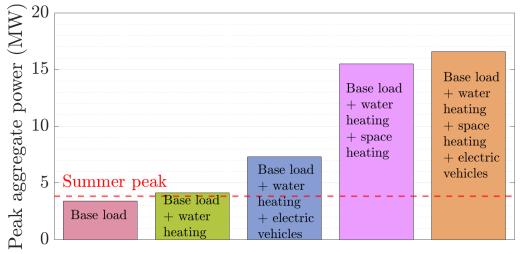

Figure 6. Space heating is the primary driver of demand peaks in the cold-weather simulations.

4. Sensitivity Analyses

4.1. Electric Vehicles

In the central case, the electric vehicles all charge on Level 2 (modeled as 240V/48A, or 11.5 kW). Would switching to Level 1 (modeled as 120V/15A, or 1.8 kW) reduce peak aggregate demand? In the week simulated here, the answer is no. While switching to Level 1 reduces the peak vehicle charging load, it also keeps vehicles charging through the night, shifting the vehicle charging peak toward the space heating demand peak around 4 AM on January 21 (see Fig. 5). These effects more or less cancel each other out, leading to negligible reduction of the weekly peak. However, switching to Level 1 charging does reduce daily peaks on other days, typically by around 10%.

As discussed in Sec. 2.2, EDGIE models the dependence of electric vehicle driving efficiencies on the outdoor temperature. Driving efficiencies typically drop off in cold weather due to cabin heating, increased energy use by battery thermal management systems,

or altered battery chemistry. Testbeds that neglect these effects risk underestimating winder demand peaks. Re-simulating the central case with constant driving efficiencies, rather than temperature-dependent ones, underestimates the neighborhood's winter demand peak by 1.3 MW (8%).

4.2. Water Heating

In the central case, all homes use resistance water heaters. Alternatively, homes could use heat-pump water heaters or hybrid heat-pump/resistance units. Heat-pump water heaters can have COPs of 2.5 or above, and therefore use substantially less power than resistance units (COP $\approx$ 1). Re-simulating the central case with all heat-pump water heaters reduces peak demand for water heating substantially, but only reduces the overall peak by about 1%. This is because even in the central case with all resistance units, water heating only contributes about 3% to the overall peak.

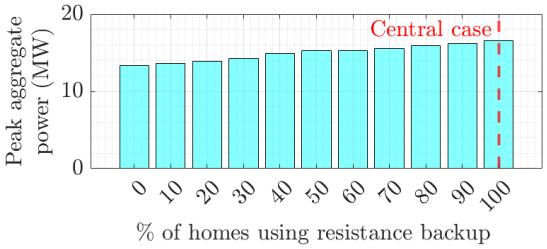

Figure 7. Using bigger heat pumps without resistance backup reduces the peak by 3.5 MW (21%).

4.3. Space Heating

In the central case, homes use air-to-air heat pumps sized for cooling. When heating load exceeds heat pump capacity, resistance backup provides supplementary heat. Alternatively, homes could use larger heat pumps sized to meet peak heating load without resistance backup (modeled here as 7 kW electric capacity, rather than 4.5 kW if sized for cooling). Fig. 7 shows the electrified neighborhood's peak demand with varying numbers of homes using resistance backup. The far right bar represents the central case. As the prevalence of resistance backup decreases, the peak also decreases, reaching 13.2 MW (24% reduction, far left bar) if no homes use resistance backup.

In the central case, users keep their indoor temperature set-points constant, night and day. Alternatively, users could reduce their temperature set-points overnight. This would save energy, as heating load scales with the temperature difference between the

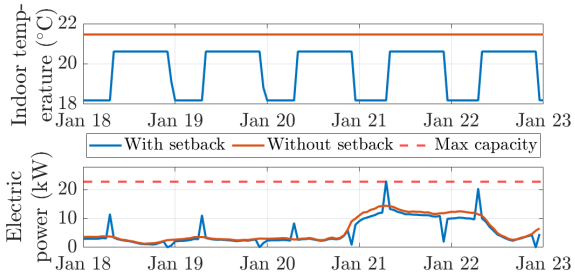

Figure 8. Warming a home back up after a night set-point reduction causes a high morning peak.

indoor and outdoor air. However, warming a home back up after a night set-point reduction can cause a new morning peak. Fig. 8 shows this effect for two similar homes: one with a constant temperature set-point (solid orange curves) and one with a lower set-point overnight (blue curve). Constant indoor temperatures keep the power used for heating (bottom plot) fairly smooth. Overnight set-point reductions, however, cause power spikes during the morning warm-up periods, sometimes reaching the combined capacity of the heat pump and resistance backup (dashed red line).

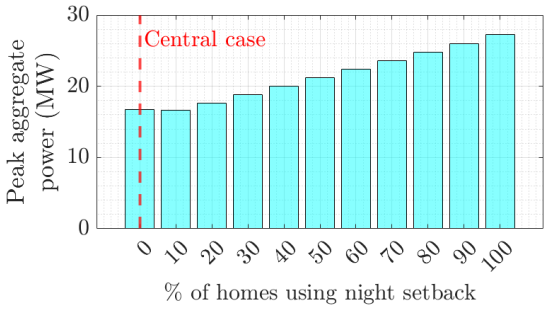

Figure 9. Lowering all temperature set-points overnight increases the peak by 9.9 MW (66%).

Fig. 9 shows the effect of night set-point reductions at the neighborhood scale. In these simulations, users raise their temperature set-points back up between 7 and 10 AM. As more homes reduce their set-points overnight, the neighborhood peak increases from 16.6 MW in the central case (far left) to 27.5 MW when every user reduces their set-point overnight (far right). This increases the neighborhood peak by 8.9 MW (66%).

Fig. 10 shows the influence on the neighborhood peak of the overall thermal resistance R between the indoor and outdoor air. This parameter models the combined effect of a home's size (hence its external surface area for heat transfer), insulation R-values, air-sealing, window sizes and U-values, *etc.* The peak increases sharply as the neighborhood-average of R

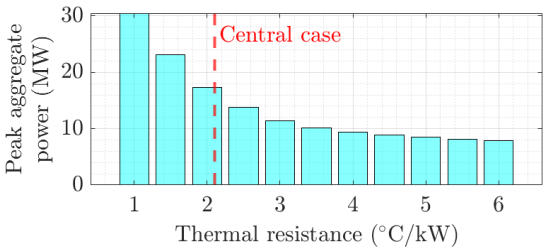

Figure 10. Increasing the mean thermal resistance to 4 °C/kW reduces the peak by 6.8 MW (41%).

drops below its value in the central case (about 2.1 °C/kW). The peak drops more slowly as R increases above the central case, with diminishing returns beyond about 4 °C/kW. The diminishing returns follow from the fact that heating load scales like $1/R$. Increasing the neighborhood-average of R to 4 °C/kW reduces peak demand by about 6.8 MW (41%).

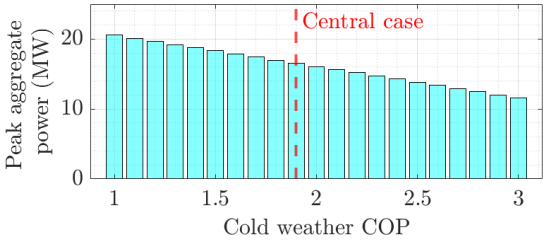

Figure 11. Increasing the mean cold-weather COP from 1.8 to 3 reduces the peak by 4.9 MW (30%).

Fig. 11 shows the influence on the neighborhood peak of the heat pumps' cold-weather COPs. These simulations follow the same set-up as the central case, including the linear dependence of COPs on the outdoor temperature. The warm-weather COP points remain at their central values (COP 3.5–4 at 7 °C). However, the cold-weather COP points sweep from their central values (COP 1.5–2 at -15 °C) down to 1 and up to 3. Increasing the neighborhood-average cold-weather COP to 3 reduces the peak by 4.9 MW (30%).

In Fig. 11, the peak decreases linearly with the cold-weather COP. This linearity may seem counter-intuitive, as a heat pump's power input typically scales reciprocally with the COP, similar to the $1/R$ dependence in Fig. 10. The linearity in Fig. 11 follows from the fact that the overall peak in the central case happens during the coldest hours, when each heat pump runs at maximum capacity and resistance backup meets the remaining load. In these hours, $\underline{\eta}\overline{p}_h + p_r = q$, where $\underline{\eta}$ is the cold-weather COP, $\overline{p}_h$ (kW) is the heat pump's electric power capacity (4.5 kW in the central case), p_r (kW) is the resistance heater power, and q (kW) is the home's heating load. Solving this equation for p_r shows that the combined power input for heating, $\overline{p}_h + p_r = \overline{p}_h + q - \underline{\eta}\overline{p}_h$, scales linearly with $\underline{\eta}$. The slope, $-\overline{p}_h = -4.5$ kW, agrees with Fig. 11.

5. Optimization Example

This section demonstrates EDGIE's optimization capabilities through the example of coordinating bidirectional electric vehicle charging with flexible space and water heating. This example has about 1.2 million decision variables: the vehicle charging and discharging powers, battery energy levels, water heating powers, water temperatures, space heating powers, and indoor air temperatures. The optimization decides these variables at each one-hour time step over five days and for each of the 990 homes and 1,980 vehicles in the neighborhood. The cost function penalizes peak aggregate demand, electrical energy use, and deviations of the indoor air and hot water temperatures from their user-specified set-points. The models (1), (3) and (8) act as constraints. The optimization also constrains the air and water temperatures to bands around their set-points. EDGIE solves the problem using CVX (Grant and Boyd (2014)), a convex optimization toolbox. Because the optimization problem is convex, CVX returns a globally optimal solution within a specified numerical precision. Optimizing the 1.2 million variables takes about 40 minutes on a 2.9 GHz processor with 16 GB of RAM.

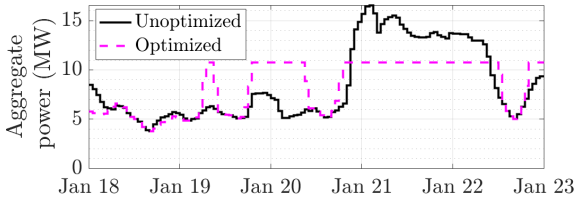

Figure 12. Optimal bidirectional charging and flexible heating reduce peak demand by 5.8 MW (35%).

Fig. 12 shows the aggregate demand in the central case (black curve) and with optimal device coordination (dashed magenta curve). Optimization reduces the peak by 5.8 MW (35%), primarily by charging the vehicle batteries off-site before the coldest nights and discharging them to power the heating systems. The indoor air and hot water temperatures also drop slightly during the coldest hours, further reducing the peak.

Fig. 13 shows the per-unit voltages in the IEEE-33 bus network in the central (black) and optimized (magenta) cases during the peak hour. In the central case, voltages drop catastrophically, reaching about 0.05 per-unit at some buses. (Ideally, per-unit voltages would remain near unity.) Voltages are higher in the optimized

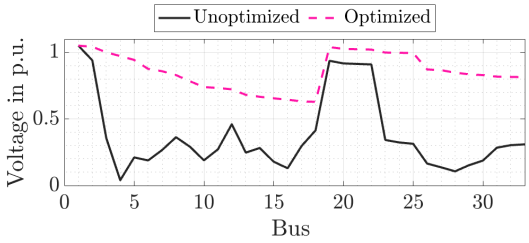

Figure 13. Without optimization, voltages collapse catastrophically. Voltages improve with load optimization, but still sag at remote buses.

case, but still drop to unacceptable levels at the furthest buses from the substation (buses #14–18; see Fig. 4 for the network topology).

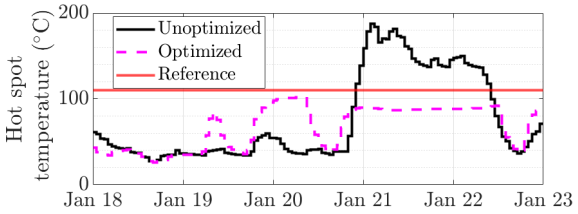

Figure 14. Load optimization keeps the substation transformer cooler, extending its life.

Fig. 14 show shows the hot spot temperature during the simulated week. Although the optimization problem does not explicitly penalize transformer degradation, peak demand serves as a useful stand-in. While the hot spot temperature reaches 190 °C in the central case, it remains below the 110 °C reference in the optimized case. This radically reduces transformer degradation, which scales exponentially with the hot spot temperature (see Eq. (12)). Table 3 summarizes the lost hours of transformer life in the two cases. For reference, transformers typically have lifetimes of 20–30 years (175–263 thousand hours).

Table 3. Transformer Loss of Life

Unoptimized (h)	Optimized (h)
4,277	9

6. Conclusion

This paper introduced EDGIE, an open-source, experimentally-validated testbed aimed at characterizing and mitigating the impacts of electrification on distribution grids. EDGIE uses real data on weather conditions, base electricity load, driving habits, home sizes, domestic hot water use, and equipment characteristics. This paper demonstrated EDGIE's capabilities by simulating vehicle charging, space and water heating, transformer degradation, and network power flows for an electrified neighborhood in a cold climate. Sensitivity analyses revealed key drivers of demand peaks. This paper demonstrated EDGIE's optimization capabilities by coordinating bidirectional vehicle charging with flexible space and water heating.

While EDGIE currently includes the main drivers of grid impacts from electrification, it does not include other equipment of potential interest. Future versions may model power line thermal dynamics, rooftop solar panels, stationary batteries, thermal storage, or backup heat sources other than electric resistance. The authors hope that EDGIE will enable researchers to investigate the grid impacts of electrification in many locations and scenarios, and to develop strategies to mitigate those impacts. These research efforts could help utilities and regulators plan grid upgrades more effectively, either by avoiding unnecessary upgrades or by sizing new power lines and transformers to more precisely match load growth. Planning grid upgrades more effectively could help keep electricity reliable and affordable as people increasingly electrify their buildings and vehicles.

7. Data Availability Statement

The data and code that support this study are available at https://github.com/priyada7/EDGIE.

References

Blonsky, M., Nagarajan, A., Ghosh, S., McKenna, K., Veda, S., & Kroposki, B. (2019). Potential impacts of transportation and building electrification on the grid: A review of electrification projections and their effects on grid infrastructure, operation, and planning. *Current sustainable/renewable energy reports*, *6*, 169–176.

Burch, J., & Thornton, J. (2012). *Realistic hot water draw specification for rating solar water heaters* (tech. rep.). National Renewable Energy Laboratory (NREL), Golden, CO, USA.

DOT. (2017). *United States Department of Transportation Federal Highway Administration: National Household Travel Survey*. Retrieved 2022, from https://nhts.ornl.gov/publications.shtml

Elmallah, S., Brockway, A. M., & Callaway, D. (2022). Can distribution grid infrastructure accommodate residential electrification and electric vehicle adoption in Northern California? *Environmental Research: Infrastructure and Sustainability*, *2*(4), 045005.

Elmoudi, A., Lehtonen, M., & Nordman, H. (2006). Effect of harmonics on transformers loss of life. *Conference Record of the 2006 IEEE International Symposium on Electrical Insulation*, 408–411.

EPA. (2021). *United States Environmental Protection Agency: Sources of Greenhouse Gas Emissions*. https://www.epa.gov/ghgemissions/sources-greenhouse-gas-emissions (accessed: 06.14.2023)

EPRI. (2018). *U.S. National Electrification Assessment* (tech. rep.). Electric Power Research Institute.
Grant, M., & Boyd, S. (2014). CVX: Matlab Software for Disciplined Convex Programming, version 2.1.
Hendron, B., Burch, J., & Barker, G. (2010). *Tool for generating realistic residential hot water event schedules* (tech. rep.). National Renewable Energy Laboratory (NREL), Golden, CO, USA.
Kircher, K. J., Cai, Y., Norford, L. K., & Leeb, S. B. (2021). Controlling big, diverse, nonlinear load aggregations for grid services by adjusting device setpoints. *2021 60th IEEE Conference on Decision and Control (CDC)*, 6377–6384.
Meinrenken, C. J., Rauschkolb, N., Abrol, S., Chakrabarty, T., Decalf, V. C., Hidey, C., McKeown, K., Mehmani, A., Modi, V., & Culligan, P. J. (2020). MFRED, 10 second interval real and reactive power for groups of 390 us apartments of varying size and vintage. *Scientific Data*, *7*(1), 375.
Novac, B., Smith, I., Wang, M., & Senior, P. (2013). Tesla-charged blumlein high-power generator. *2013 19th IEEE Pulsed Power Conference (PPC)*, 1–6.
NREL. (2023). *Distribution grid integration unit cost database* (tech. rep.). National Renewable Energy Laboratory (NREL), Golden, CO, USA.
Swift, G., Zocholl, E., Bajpai, M., Burger, J., Castro, C., Chano, S., Cobelo, F., De Sa, P., Fennell, E., & Gilbert, J. (2001). Adaptive transformer thermal overload protection. *IEEE Transactions on Power Delivery*, *16*(4), 516–521.
Tarroja, B., Chiang, F., AghaKouchak, A., Samuelsen, S., Raghavan, S. V., Wei, M., Sun, K., & Hong, T. (2018). Translating climate change and heating system electrification impacts on building energy use to future greenhouse gas emissions and electric grid capacity requirements in California. *Applied Energy*, *225*, 522–534.
White, P. R., Rhodes, J. D., Wilson, E. J., & Webber, M. E. (2021). Quantifying the impact of residential space heating electrification on the Texas electric grid. *Applied Energy*, *298*, 117113.
Yuksel, T., & Michalek, J. J. (2015). Effects of regional temperature on electric vehicle efficiency, range, and emissions in the United States. *Environmental science & technology*, *49*(6), 3974–3980.
Zhou, E., & Mai, T. (2021). *Electrification futures study: Operational analysis of us power systems with increased electrification and demand-side flexibility* (tech. rep.). National Renewable Energy Lab.(NREL), Golden, CO (United States).
Zimmerman, R. D., & Murillo-Sanchez, C. E. (2020). Matpower (version 7.1) [software].