Protecting residential electrical infrastructure through advanced control: The first field results

Elias N. PERGANTIS[1*], Levi D. REYES PREMER[1], Alex H. LEE[1], PRIYADARSHAN[1], Haotian LIU[1], Eckhard A. GROLL[1], Davide ZIVIANI[1], Kevin J. KIRCHER[1*]

[1]Ray W. Herrick Laboratories, School of Mechanical Engineering, Purdue University
West Lafayette, 47907-2099, USA

*Corresponding Author: epergant@purdue.edu, kircher@purdue.edu

ABSTRACT

Electrifying vehicles and residential appliances can significantly reduce greenhouse gas emissions in areas with clean electricity. However, the electrical infrastructure in most older houses, was not designed to accommodate peak current draws from large loads, such as space heating, water heating, and vehicle charging. Upgrading a house's circuit breaker panel and/or electrical service (the wires that connect a home to the utility) can be a significant barrier to rapid electrification. This study develops a novel two-level control architecture that robustly maintains the total current draw in an all-electric home within the safe limits of existing panels and service. The control system adjusts the temperature set-points of an electric resistance water heater and an air-source heat pump. The high-level controller adjusts set-points over a receding forecast horizon, while the low-level controller monitors the real-time conditions and switches off appliances if necessary. The system is tested over 31 days in an occupied, 208 m^2 house in a cold climate. Test conditions include outdoor temperatures as low as -20 ℃. The controller successfully maintains the whole-home current within the safe limits of electrical panels and service rated at 100 A, a common rating for older US homes. The potential value of this work is to allow older homes to safely electrify without upgrading electrical panels or service. If electrical codes permit, this could save a typical homeowner on the order of two to ten thousand dollars and eliminate significant delays.

1. INTRODUCTION

The deep electrification of buildings to move away from burning fossil fuels, including water heating, space heating, and other appliances such as cooktops, ovens, dryers, and washers can lead to significant carbon emission reductions when run on low-carbon electricity (IEA, 2023). However, a significant barrier to these efforts lies in the capacity of the electrical infrastructure both on the building side (e.g., breakers, wires, etc.) and on the distribution side (Sharma et al., 2020). There is a mismatch between the existing limits of electric infrastructure and the increased demand from electrified residential buildings. This study focuses on the mitigation of a bottleneck to electrification in residential buildings, namely the capacity of circuit breaker panels and the electrical service that connects the house to the nearest distribution transformer. It is common for residential buildings to have legacy breakers, sized at the amperage required to meet a few plug loads but not major appliances such as electric water heaters or heat pumps (NEC, 2014). As more appliances shift from fossil fuels to electricity, the existing breakers often require an upgrade. This is a significant cost for homeowners, requiring an electrician to switch off power, disconnect and reconnect the wires between the different appliances and the main utility distribution line and to install the new breakers for the house. This whole process can cost $2,000 to $10,000 for a household in the U.S., with the higher end of the range corresponding to values that also require significant upgrades to the wiring connecting to the main distribution (Lindsey, 2023). Particularly burdened by this are homes in cold climates, which often require high-power resistive backup heat elements to meet heat demand on the coldest days.

To address this issue and provide a practical pathway for electrification, this paper formulates and addresses a new research question: *If a main breaker has adequate sub-breaker spaces or additional wiring to accommodate the further electrification of the residential building, can smart coordination of the appliances meet the house's needs without requiring an upgrade?* In this study, a novel multi-layer controller focusing on current limiting was developed,

followed by an experimental demonstration of this technology in a fully electrified house in a cold climate in the northern US. The experimental testbed, the DC Nanogrid House at Purdue University, would nominally require at least 200 A of panel and service capacity under today's electrical standards. However, the controller developed here successfully ran the house in the safe region of a 100 A panel over a period of 30 days during the peak heating season. As far as the authors are aware, although there have been related demonstrations of demand-charge reduction in commercial buildings (Kim et al., 2015, Kim and Braun, 2018, Zhang et al., 2022) and simulations of demand-charge reduction in residential buildings (Paterakis et al., 2015), the unique problem of robustly constraining total current, rather than reducing an economic penalty on peak power demand, has been little discussed and never before tested in the field (Khabbazi et al. 2024).

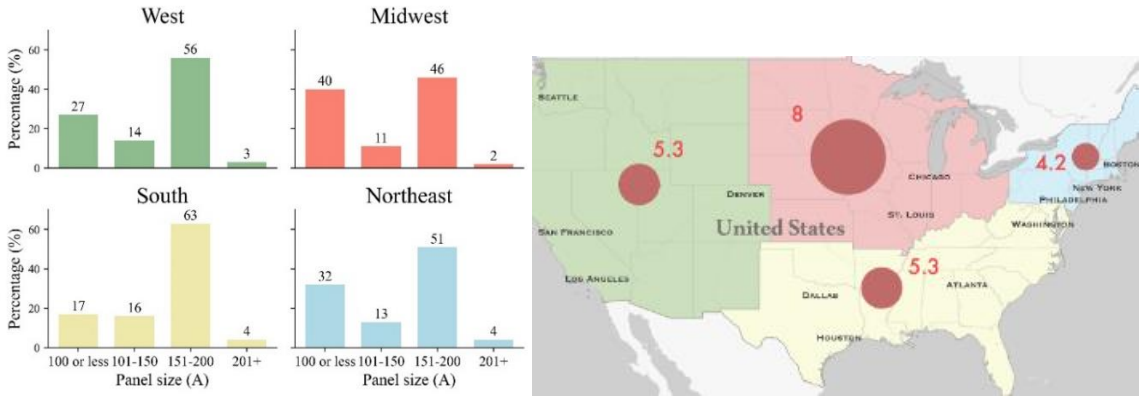

Figure 1: (Left) Percentage breakdown of panel sizes by four census regions. The West, South, and Northeast census regions have the majority of panels sized between 151 amps and 200 amps. The Midwest region has the majority of panel sized 150 amps or below. **(Right)** Number of panels (in millions) that require upgrade in the case of full electrification with back up heat and electric vehicle adoption.

2. UNDERSTANDING THE PROBLEM

2.1 Circuit breaker sizes in the US

To truly understand the impact of electrification and adoption of electric vehicles, and the necessary breaker panel upgrades to accommodate such upgrades, we first need to understand the current state of electrical infrastructure for residential buildings. While the problem of full electrification is relevant worldwide, our scope is limited to the US. To the best of the authors' knowledge, there exists little to no national-level data on panel sizes, and only a couple of actual surveys were performed investigating the panel sizes in residential homes. First, Pecan Street (Merski, 2021) surveyed 263 residential homes in Austin, Texas. They noted that the square footage of the house, the year the house was built or the year that it was last renovated, and the fuel source of home appliances have the most impact on the rated size of breaker panels. However, these only serve as an approximation for the actual panel size and much of this data is not publicly available (e.g., the last renovation date) or it varies greatly from place to place (e.g., local fuel infrastructure and fuel pricing). They also noted that homes built with combustion fuel assistance such as natural gas heating will often have electrical panels rated below 200A. The 2023 EIA survey data (EIA, 2023) reports that the Midwest region had the highest annual natural gas and propane use in residential households, both in total and on average per household. This finding echoes that of EPRI's recent report (Lindsey, 2023) that surveyed 1,857 homes across the US and categorized into the four census regions: West, South, Midwest, and Northeast. The results of this study are highlighted in Fig. 1. It can be seen that under the electrification scenario described in Lindsey 2023, a significant number of households are expected to be required to upgrade their electrical panels especially in cold climates, with a total cost in the several billions USD (EIA, 2023).

2.2 Safety limits and the DC House

Panel size upgrades increase the rated capacity to allow safe operation and prevent the main circuit breaker in a home from repeatedly tripping which can shorten its lifetime. The circuit breaker in a home is an essential device that prevents fires when the current approaches unsafe operation under the designed current load. To prevent fire hazards,

the breaker turns off, in other words trips, preventing electricity to be distributed within the home or the individual line the breaker is attached to. A breaker usually has two tripping mechanisms, which are shown in Fig. 2. The first one is a thermal tripping mechanism that initiates from heat within the breaker. As the current increases, the temperature of the breaker will rise, and the breaker will trip when it is approaching a hazardous state. This can happen after minutes, or seconds depending on the amount the current exceeding the rating. The greater the current, the faster the breaker will trip. The second mechanism is electromagnetic tripping, which happens when current multiples much greater than the rated current. The tripping will happen with no intentional delay, on the order of seconds to milliseconds for safety precautions (Larsen, 2008). Both mechanisms can be identified in the trip curve developed by the manufacturer, which provides the range of times the breaker will trip at a certain multiplicity of the rated current.

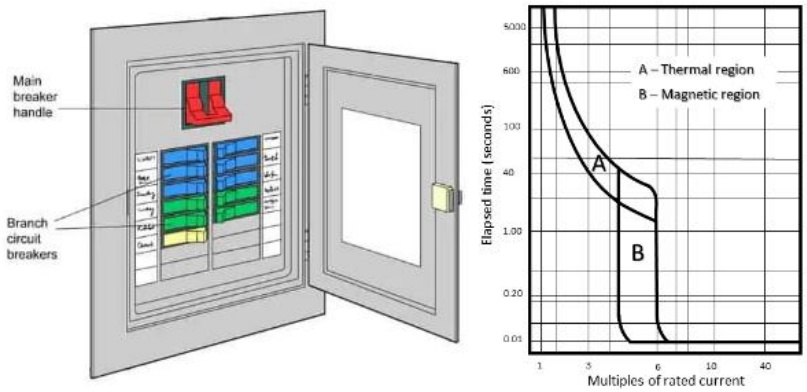

Figure 2: (Left) An example home electrical panel with the main breaker and sub breakers to the left (Alion 2021) **(Right)** Example breaker time-current curve (trip curve). The possible failure region within the DC home resides in the thermal tripping region.

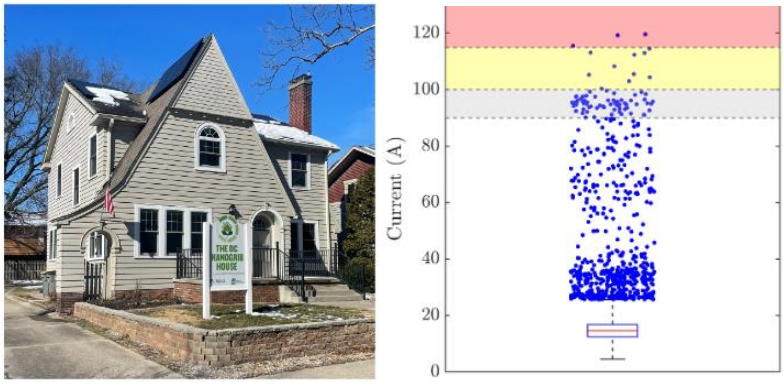

Figure 3: (Left) The fully electrified test house (208 m^2). **(Right)** Historical (prior to smart controller) current draw at the site (box and whisker plot). All data points are 5-minutely. In grey is the undesirable region, yellow is above 100 A but still technically safe (before thermal trip region), and red is when there is a chance the breaker will trip.

In a panel there is the main circuit breaker that has sub breakers that monitor either individual large appliances or areas with smaller loads (e.g. bedrooms). For this study, we assume the sub-breakers are correctly sized and that the main failure region is only within the main circuit breaker. The main circuit breaker will therefore only be referred to in the latter sections. The field demonstrations in this paper took place in the DC Nanogrid House, pictured in Fig. 3. This is a fully electrified home on the Purdue campus, in West Lafayette IN. The reader is directed to the following references for more information on the experimental setup (Pergantis et al. 2024). The house is conditioned via a central air-to-air heat pump with backup resistance heating elements. The maximum rated electrical compressor power is 4 kW, and supplemental heating elements in three stages (9.6, 14.5, 19.2 kW). A 50-gallon electric water heater provides hot water, with a rated heating resistance of 4.5 kW. Using the National Electric Code (NEC) article 220 (NEC, 2014), the electric panel in the DC house is rated at 200 A. Therefore, the safety limit of a 100 A breaker is emulated as was in place at the house before the electrical upgrades took place.

2.3 Control problem

To first identify the baseline operation, we analyzed the total current draw in the house during the 32 days in the winter of 2022-2023 period. This is shown in Fig. 3 (right). Based on the trip curve of typical 100 A breaker, the house had potentially three points of failure over the 32-day period seen in the red region. Depending on the age of the breaker and the breaker manufacturer, since all breakers do not have the same exact time current curve, the points close to the red region could become a part of the tripping region. The number of occupants largely affected these current peaks in the data as well. The more occupants, the more likely there will be synchronization of equipment (e.g. water heater and heat pump). In the period of December and January of 2022-2023, there were only two occupants, one of which was on vacation for part of December. For December and January of 2023-2024, during our experimental period, there were three occupants, one of which was on vacation for part of December. To quantify this change in occupancy, we compared the home's hot water usage between the winters of 2022-2023 and 2023-2024 in Fig. 4. Hot water usage was normalized over the house's average shower of 12 minutes, consuming roughly 20 gallons of hot water. As seen in Fig. 4, hot water usage was often three times the previous year's winter's usage. This results in a greater likelihood of synchronization of the water heater with other equipment, a more difficult problem for the 2023-2024 Winter.

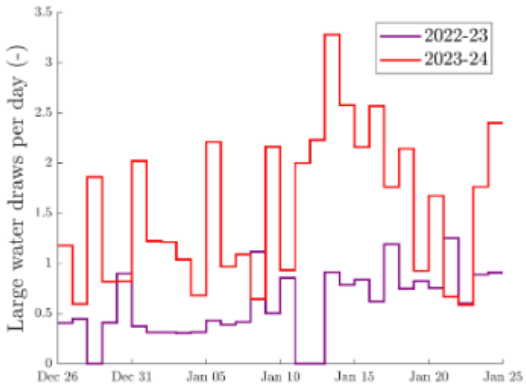

Figure 4: Hot water usage compared between 2022-2023's and 2023-2024's Winter. In 2022-2023's Winter, there were only two occupants while there were three in the 2023-2024 period.

Critically, failure in the form of breaker trips occurs when large resistive loads from the electric water heater and the heat pump backup stages are synchronized. This can occur either during normal operation or during defrost. The presented controller in Sections 3 and 4 has two separate layers to achieve a disaggregation of resistive loads in the house. The slow layer (5-minutely interval) acts as a supervisory planner that aims to improve occupant comfort and disaggregate space heating and water heating loads. The fast layer (30 second interval) has both reactive features, i.e., what is the real-time current at the house, as well as some predictive features based on defrost forecasts as well as predicted load inputs from the fast layer.

3. HIGH-LEVEL CONTROLLER

3.1 Envelope model

The slow layer uses a supervisory model predictive control (MPC) formulation. As per Pergantis et al. (2024), a thermal resistance capacitance (RC) 2R1C is identified using measurements from the house (shown in Fig. 5). Unlike previous work (Kim and Braun, 2022), that use a Kalman filter to identify hidden states, a simple RC model with observable states is found to perform well (Pergantis et al. 2024b). This has been seen in previous simulations for detached homes with air-to-air systems (Blum et al., 2019). Training must be performed due to the difference in envelope dynamics when the time steps are reduced from hourly to 5 minutes. Unlike the methodology in Pergantis et al. (2024) where a static model was used, a dynamic machine learning model is used to predict the so-called heat supply disturbances not captured by the RC model, such as transient and non-linear envelope effects, infiltration, solar gains, people activity and others. A dynamic term is used by incorporating historical on-site observations and serves as a catch-all error term.

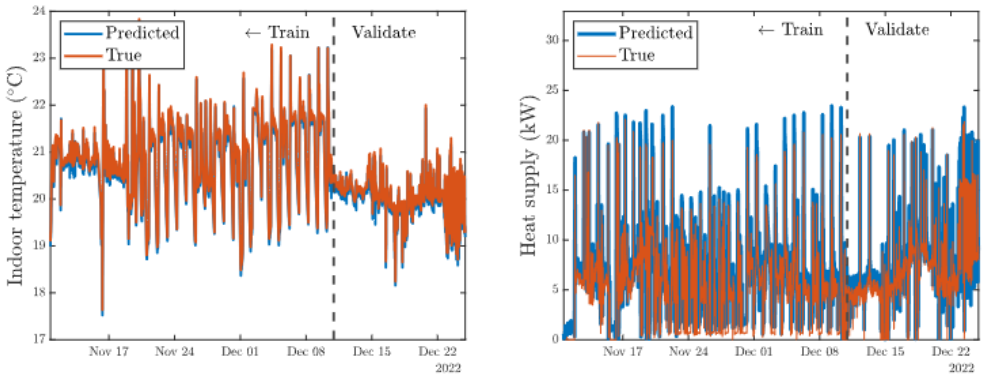

Figure 5: (Left) The full model's step-ahead predictions (blue curves) of the indoor temperature and the HVAC thermal power **(right)** match the targets (orange curves) well in training (left two-thirds of data) and validation (right one-third). The validation RMSE temperature and power prediction errors were 0.2 ℃ and 1.2 kW, respectively.

3.2 Equipment models

For the heat pump, the coefficient of performance is obtained as quadratic function of the outdoor temperature. Dropping the indoor temperature dependance of the coefficient of performance allows for a linear formulation. The current draw is modeled as a sum of compressor current draw plus that used by the backup resistance. A power factor of 0.8 (constant) is used for the compressor, while one is used for resistive elements since they provide Joule heating. For the water heater, a mixed-tank single-node model was used, with the initial temperature of the tank being a mixed sum of temperatures read by a thermocouple at the top of the tank and another at the bottom (75% upper, 25% lower as per Reyes Premer et al. 2024). A k-nearest neighbors' algorithm is used to predict future water use based on a historical database of the last 30 days, with the last 8 hours being the prediction feature. The days are updated to capture occupancy changes.

3.3 Rest of home & robust formulation

So far, models for the water heater, heat pump and envelope have been presented. The supervisory controller requires knowledge of the "rest of home" current draw, i.e., everything which is not controllable, sockets, appliances, etc. In this study, unlike the power limiting work of Paterakis et al. (2015), full freedom is given to the occupants to operate their appliances, and the optimization works around these. A robust formulation is adopted (Rakovic, 2019), where the 99th quantile of peak current draw is considered, binned in hourly time intervals over which a zero-order hold is performed (Fig. 6). For the robust formulation, a sweep of water forecasts is obtained, while for the envelope, a 95th percentile of worse heat loss disturbance.

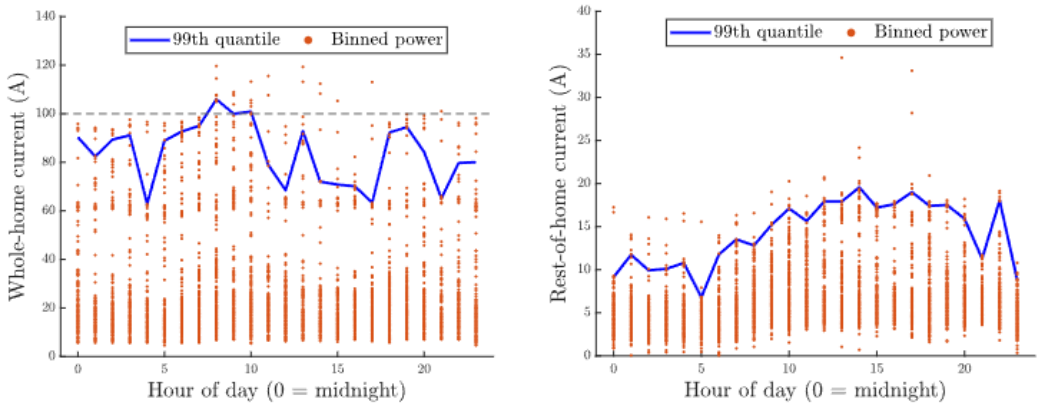

Figure 6: (Left) Whole building binned current draw including an HP, backup heat, an EWH, oven, dishwasher, dryer, washer, and sockets. **(Right)** Rest of home not including the HP, backup heat, and EWH.

3.4. Optimal controller

The optimization is posed in CVX and solved using Gurobi 9 (Grant and Boyd, 2014). Due to the on-off resistive loads, and for the purpose of capturing the true current peaks and to coordinate the electrical resistive elements of the water heater and backup for the heat pump, the problem is posed as a mixed integer linear problem (MILP). At each time instance, as per Fig 7. (right), forecasts of the ambient weather conditions are obtained using a weather service (Oikolab) to predict the coefficient of performance of the heat pump, the envelope disturbance heat loss (alongside historical values), as well as the predicted water use (not weather dependent). Real-time measurements of the indoor temperature and water temperature are extracted from the water heater and heat pump using their thermostats' API calls and stored on a remote server (InfluxDB hosted on Digital Ocean, Pergantis et al., 2024). Constraints are applied to the minimum allowable temperatures (air/water), as well as capacity limiting of the heat pump based on the coefficient of performance mapping. Due to the fast time scales (5 minutes), a thermostat model is also utilized to capture indoor zone tracking dynamics, omitted here for brevity. In the following expressions, S is the number of robust forecasts (superscript s indicating one of them). It was found that eight scenarios had good performance without sacrificing computation resources. The time step Δt was 5 minutes, while the forecast horizon K was 12 hours. The subscript k denotes a discrete time instance between 0 (now) and K. The following objectives were used:

- **Indoor temperature and water comfort tracking:** It was found that a tracking preference ensures that enough capacity is stored both in the indoor air and water heater, that allows for a better performance of the controller while improving occupants' comfort. This is given in Eq. 1, where π_t is weight to air and π_w is the weight to water set-point tracking. These terms were tuned and at the start of the optimization. The tilde notation denotes a set-point, *pref* to the set-point preference of the occupants and *w* to water states.

$$\frac{1}{S}\sum_{s=1}^{S}\sum_{k=1}^{K}\left(\pi_t\left|\tilde{T}^s_{pref,k}-\tilde{T}^s_k\right|+\pi_w\left|\tilde{T}^s_{w,pref,k}-\tilde{T}^s_{w,k}\right|\right) \tag{1}$$

- **Minimize cost of electricity:** Over the forecast period, the cost of electricity is continuously optimized under a flat utility rate in place in the location of testing (Duke Energy Residential Rates, Indiana).
- **Minimize violation of peak current:** The objective to avoid current overcharge, shown in Eq. 5, is posed as a soft constraint with a significant magnitude to its weight π_I (10^5), so that the optimization avoids violating it. The safety limit, I_{Safety} was set to 90 A (90% of 100 A panel), with the other terms being the rest-of-home current draw (I_{Home}, Section 3.3), the heat pump net current draw (I_{HP}), and the water heater current draw (I_{EWH}). The max over both the range of scenarios and over the forecast horizon is used to ensure that the objective is never violated.

$$\pi_I \max_S\left(\max_K\left(\max\left(I^s_{HP,k}+I^s_{EWH,k}+I^s_{Home,k}-I_{Safety,k},0\right)\right)\right) \tag{2}$$

Once the optimization problem is solved, a unique action for the water and air temperatures is sent across all the scenarios to the low-level controller, which will either be allowed or overruled based as discussed in Section 4. A unique action is ensured by constrained on the set-points to be the same for the next time-step over all scenarios.

4. LOW-LEVEL CONTROLLER

A higher-level optimizer operating at a slow time step can satisfy several objectives such as ensuring the thermal comfort of occupants and minimizing energy use and energy costs. For the purposes of this project, it can also help mitigate the chance of multiple appliances synchronizing and exceeding the maximum current limit by coordinating individual appliances appropriately. However, the controller cannot perfectly capture the nondeterministic pattern of occupant behavior (e.g. when someone decides to run the dryer or the oven) or the behavior of large electric appliances (e.g. the heat pump entering defrost mode and the backup resistance heat being activated). As such, a lower-level fast controller layer has been designed primarily to continuously monitor the real-time whole home conditions. Critically, the fast controller has two primary functions, shown in Fig. 7 (right):

- **Reactive real time operation:** It is possible that due to an unforeseen water draw, or other events an extreme peak might occur which was not prevented by the operation of the high-level MPC system. When that happens, the fast layer lowers the set-points for the heat pump and water heater, leading to a quick drop of the current. As such the net current draw of the house is monitored at a 30 second interval. Once the violation has passed, the low-level controller allows the set-points requested by the high-level controller.
- **Predictive operation:** The low-level controller has two types of predictive logic embedded. Firstly, MPC communicates the predicted stages of backup heat (z_{aux}, in Fig 7 right). If the baseline draw (non-HVAC) increases above a safety threshold and a predicted stage of backup is anticipated, the set-point is decreased. The second logic involves, predicting if a defrost is going to occur based on the supply temperatures. As frost builds up on the outdoor heat exchanger, the supply temperature drops, something that can be captured.

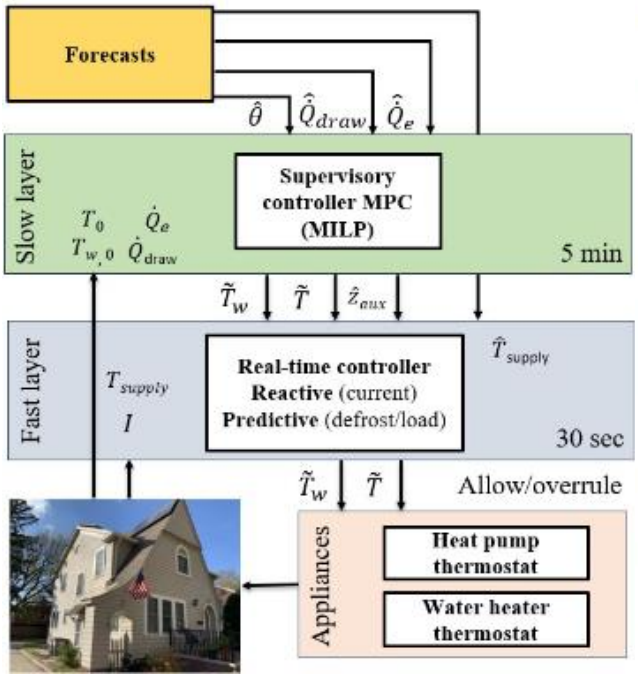


```
Fast Layer Controller:
while true do
    //Reactive
    Read whole house current;
    Read slow layer set-points for heat pump and water heater;
    Read slow layer predictions on aux use to meet load;

    if unexpected water usage and water heater is charged then
        Set water heater to 120;
    end
    if unexpected stage II or III aux heat use then
        Lower water heater set-point to 110;
    end
    if whole house current > safety limit then
        Turn water heater off;
        Lower compressor set-point;
    end

    //Proactive
    Make predictions on aux heat use for defrost;
    if Defrost is predicted then
        Turn water heater off;
        if Stage I or II defrost & Total current > safety limit then
            Lower compressor setpoint;
        else if Stage III defrost then
            Lower compressor setpoint;
        end
    end

    if Aux heat is active to meet heating load then
        Turn water heater off;
        if Stage II aux & whole house current > safety limit then
            Lower compressor setpoint;
        else if Stage III aux then
            Lower compressor setpoint;
        end
    end

    Correct over-predicted defrosts;
    If no safety measures need to be taken, push slow layer's setpoints
end
```

Figure 7: (Left) Flow chart of slow-layer MPC, with prediction denoted by the hat ^ notation and set-points by tilde ~. The fast layer can overrule or allow the request of the slow layer based on the real-time measurements. **(Right)** Flow chart of the fast controller, indicating both predictive and reactive features (Section 4).

5. EXPERIMENTAL RESULTS

5.1 Performance of controller

Experiments were run at the DC House, between December 25th 2023, and January 25th 2024, for a total of 31 days. These were compared with historical data collected over a similar period (December to January) of 2022 2023. Both test MPC days and baseline operation included a cold-snap period, with ambient temperature falling below -20 ℃ . The mean ambient temperature for the non-MPC baseline was 1.25 ℃ , while for MPC -2.7℃ . Concerning the indoor conditions, the set-point for non-MPC was 20.5 ℃ , while for MPC it was 20.7℃ ; for the water temperatures, the average supply temperatures were 47.8 ℃ (non-MPC) and 47.6 ℃ (MPC). These indicate that the occupants' preferences were met satisfactorily, something corroborated by on-site visits and questionnaires. It should be noted, however, that for MPC the testing conditions were significantly more challenging because of (1) the lower ambient temperatures and (2) the number of occupants as mentioned in Section 2. Due to last year's cold snap coinciding with the Christmas holiday break, there was limited occupancy at the house. Additionally, last year the house had two

baseline occupants, while this year it has three. While the MPC experiment was in place, the two controllers, the fast and slow layer were working in tandem to keep the current below a rated limit. The time series of the testing data is shown in Fig. 8 (left) while Fig. 8 (right) shows the comparison across two testing days. From Fig. 8, it can be seen that during non-MPC operation frequently the house surpasses that safety limit regardless of weather conditions primarily driven by defrost cycles.

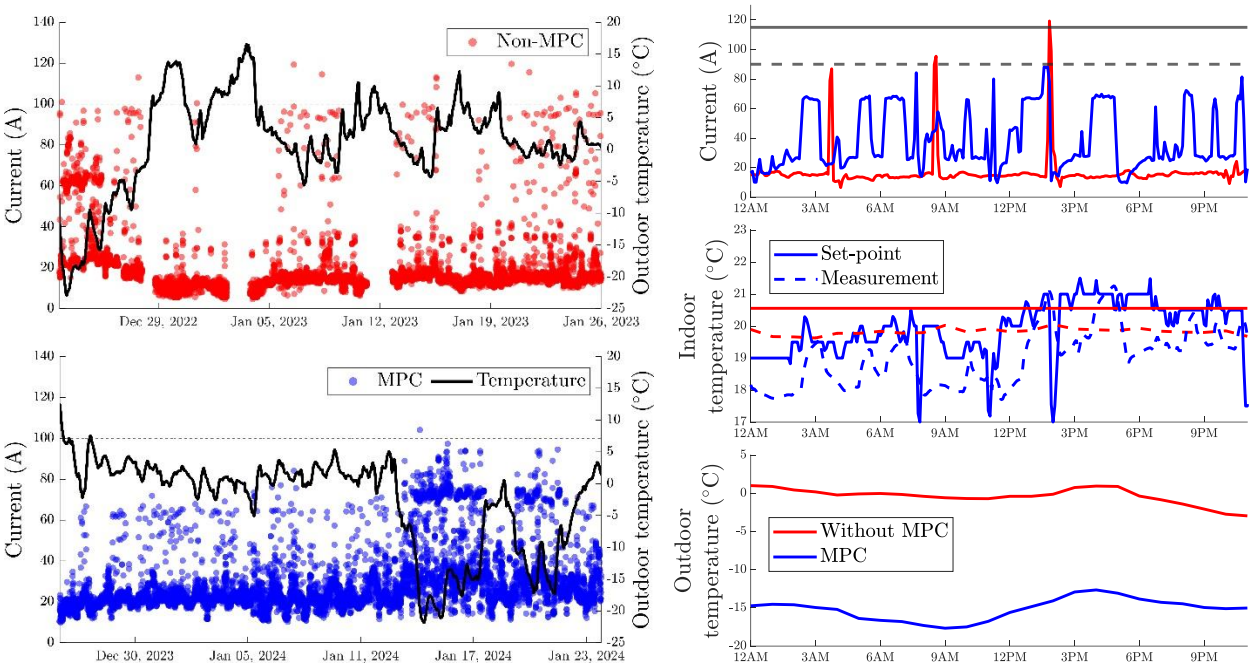

Figure 8: (Left) Net current draw across winter of 2023 and winter of 2024 (5-minutely measurements). It can be seen that for non-MPC operation, frequently the current draw exceeded 100 A. **(Right)** Comparison of two days on record, one with and one without. Although the non-MPC operation was on a milder day (-0.4 ℃ vs -15.2 ℃ for MPC), twice it surpassed the desirable operation region of the breaker (90 A) and once the potential trip (115 A). The peak due to coordination of water heater and defrost reached 120 A, while for MPC, it never surpassed 90 A.

The multilayer controller, on the other hand, was successful in maintaining the peak limit condition across both mild and cold days. The controller is successful in reducing how frequently the net current surpasses 90 A, which was 0.61 times per day for the controller as opposed to 2.2 for the baseline operation despite the more challenging testing conditions. The controller's performance is further elucidated in Fig. 8 (right). Although the baseline operation is across a mild heating day, without the lock-out temperature for the high-stage defrost, it can be seen that the non-MPC current draw frequently exceeds 100 A, often reaching values as high as 120 A. peak is in the unsafe region. A defrost that synchronized with a large water draw at around 2 PM, had a five minutely peak of 120 A. For the proposed supervisory controller, even though the temperature reached as low as -18℃ , the current limit was maintained. Firstly, performing a nighttime setback along with a slow recovery ensures that no stage II/III backup heat is used to meet the building load. Adequate preheating of the water heater and disaggregation further decouples the resistive loads in the house. Finally, at 3 instances the secondly current measurement surpassed the safety limit, the fast layer lowered the temperature set-point (sharp drops) for 5 minutes to ensure a safe operation. This was performed because as a defrost is stopped and heat pump switches to heating mode, the resistive load remains on to aid the system's fast recovery (manufacturer logic) which can result in a large peak. Lowering the set-point avoids this second peak. The optimization also uses the period of higher ambient temperature (3-6 PM) window, to max out the compressor when it's efficient to do so.

When comparing the safety performance of the current limiting controller to last year, the controller was excellent at mitigating the current below the breaker tripping zone (Fig. 9). The homes operation in the undesirable region with the controller was much less frequent as well due to a design choice to have the low-level controller mitigate the

current below 90 A. This 90 A threshold in the fast controller was partially chosen due to time delays in signal communication between the equipment and to provide a buffer period for any uncontrolled equipment turning (e.g. washer and dryer). During one 5-minute interval, the current in the house exceeded 100 A, reaching 104 A. This was due to a more extensive communication delay allowing the water heater to stay on during part of a defrost cycle. Despite this, the peak was well below the potential tripping zone. The controller was able to successfully prove safe operation is possible in an electrified without a panel upgrade in extreme weather conditions. Under a no smart control operation (Fig. 9, left), 5 potential thermal tripping instances of the breaker occurred (point in or near red region).

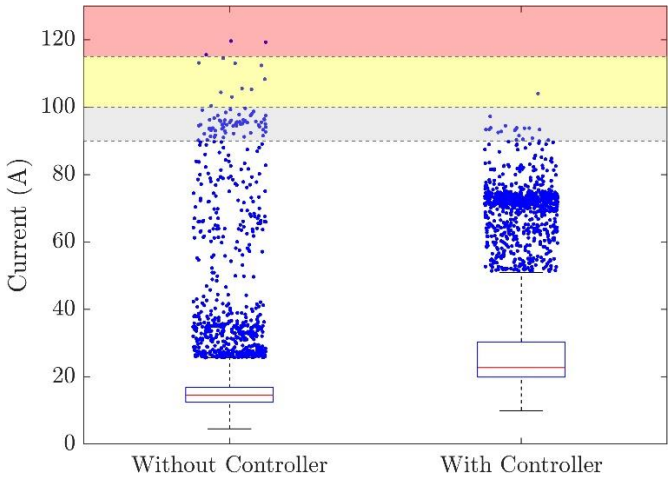

Figure 9: The current draw at the test-site (5-minutely) with and without the controller. Using MPC, we stayed within the safe region of the breaker's operating zone and only went above 100 A only once (still safe region).

6. CONCLUSIONS

This work is focused on presenting a practical and simple predictive controller aimed at multi-device coordination in residential buildings for the purpose of peak current limiting control. With aging infrastructure, and increased penetration of electrical appliances in previously primarily gas-operated homes, this is expected to be a significant bottleneck to electrification with an expected of cost of $2,000-10,000 per household, proving a critical barrier to electrification for low-income households. Recent studies have highlighted that this will impact millions of households in the US in the coming years. This study presents the following:

- A novel two-layer control predictive controller formulation is presented. The high-level controller coordinates two major appliances, an electric water heater, and an air-source heat pump, based on forecasts of future weather, disturbances, and occupancy patterns. The set-points of these appliances are adjusted such that the peak current draw of the home is disaggregated. A low-level controller monitors the real-time conditions for potential disturbances not captured by the supervisory layer. The combination of these controllers achieves both comfort, as well as protection of the electrical infrastructure.
- An experimental demonstration of peak current limiting control in a real-world residential building is performed. This is a fully electrified home (208 m^2) in a cold-climate location with temperatures during testing as low as -22 °C. The test site has an upgraded panel capacity of 200 A, while before the electrification project, it had 100 A. It was shown that with smart controls the peak current draw of the home was kept in the safe region of the 100 A panel for over a month of testing. If instead, the panel had not been upgraded and with no smart controls in place, 5 breaker trips might have occurred, leaving the home without power during an extremely cold period.

Future work will be placed on making the controller more scalable, by lowering the engineering and deployment costs. Additionally, efforts are made to extend this methodology to multifamily homes as well as, operation with on-site electric vehicle charging. The major barrier to this kind of engineering solution are the electrical codes in place which as of now (NEC, 2014) do not allow for control solutions. Further work is ongoing to investigate mathematical guarantees on the ability of the controller to protect the panel, and investigate its limitations across building types, equipment capacities and climates. This technology has the potential to help millions of homes to reduce their electrification cost from a bottom-up approach and accelerate the reduction of gas-powered appliances.

REFERENCES

Delmastro, C., Chen, O. (July 2023). Buildings. Tech. rep., International Energy Agency.

Sharma, N., Acharya, A., Jacob, I., Yamujala, S., Gupta, V., Bhakar, R. (2021). Major blackouts of the decade: Underlying causes, recommendations and arising challenges. In: 2021 9th IEEE Intern. Conf. on Power Systems (ICPS), IEEE, pp. 1–6.

NFPA 70: National Elecrical Code (NEC), Article 220, 2014.

Paterakis, N. G., Erdinç, O., Bakirtzis, A. G., Catalão, J. P. S. (2015). Optimal household appliances scheduling under day-ahead pricing and load-shaping demand response strategies. IEEE Trans. on Ind. Informatics, 11(6).

Merski, C. (2021). Addressing an electrification roadblock: Residential electrical panel capacity.

EIA. (2023). Residential energy consumption survey. Tech. rep.

Lindsey, D. (2023). U.S. residential electrical panels: How many need to be upgraded. Tech. rep., EPRI.

Larsen, E. (2008). A new approach to low-voltage circuit breaker short-circuit selective coordination. In: 2008 IEEE/IAS Industrial and Commercial Power Systems Technical Conference, pp. 1–7.

Pergantis, E. N., Priyadarshan, Theeb, N. A., Dhillon, P., Ore, J. P., Ziviani, D., Groll, Kircher, K. J. (2024). Field demonstration of predictive heating control for an all-electric house in a cold climate. Applied Energy, 360, 122820.

B. D. Less, N. Casquero-Modrego, I. S. Walker, Home energy upgardes as a pathway to home debcarbonization in the US: A literature review, Energies (2022). Residential energy consumption survey (recs). (2020).

Zhang, K., Prakash, A., Paul, L., Blum, D., Alstone, P., Zoellick, J., Brown, R., Pritoni, M. (2022). Model predictive control for demand flexibility: Real-world operation of a commercial building with photovoltaic and battery systems. Advances in Applied Energy, 7, 100099.

Kim, D., Braun, J., Cai, J., Fugate, D. (2015). Development and experimental demonstration of a plug-and-play multiple RTU coordination control algorithm for small/medium commercial buildings. Energy and Buildings, 107.

Kim, D., Braun, J. E. (2018). Development, implementation and performance of a model predictive controller for packaged air conditioners in small and medium-sized commercial building applications. Energy and Buildings, 178.

Blum, D., Arendt, K., Rivalin, L., Piette, M., Wetter, M., Veje, C. Practical factors of envelope model setup and their effects on the performance of model predictive control for building heating, ventilating, and air conditioning systems. Applied Energy, 236, 410–425.

Rakovic, S.V. (2019). Robust Model Predictive Control, Springer London, London, 2019, Ch. 2, pp. 1–11.

Reyes Premer, L.D., Semmelmann L., Pergantis E.N., Groll, E.A., Ziviani, D., Kircher, K.J. (2024). Field Demonstration of Predictive Heat Pump Water Heater Control, *8th Intern. High Perf. Buildings Conf.* Paper 3536.

Khabbazi, A.J., Pergantis, E.N., Reyes Premer, L.D, Lee, A.H., Ma, J., Liu, H., Henze, G.P., Kircher, K.J. (2024). What Have We Learned From Field Demonstrations of Advanced Commercial HVAC Control? 8th International High Performance Buildings Conference. Paper 3503.

Pergantis, E. N., Park, J., Bird, T.J., Ore, J. P., Ziviani, D., Kircher, K. J. (2024). Open-source tool-box for identifying envelope dynamics of detached residential buildings, 8th International High Performance Buildings Conference. Paper 3139.

Ohnaka, T. T., Watanabe, Y. (1994). The effects of variation in body temperature on the preferred water temperature and flow rate during showering. Ergonomics, 37, 541–546.

Priyadarshan, E., Pergantis, C., Crozier, K., Baker, K., Kircher, K. J. (2024). Edgie: A simulation test-bed for investigating the impacts of building and vehicle electrification on distribution grids. In: Proceedings of the Hawaii Intern. Conf. on System Sciences.

M. Grant, S. Boyd, CVX: Matlab software for disciplined convex programming, version 2.1

ACKNOWLEDGEMENTS

The Center for High-Performance Buildings (CHPB) at Purdue University supported this work (project CHPB-26-2023). E. Pergantis was also supported by the Onassis Foundation as one of its scholars, as well as the American Society of Heating and Refrigeration Engineers (ASHRAE) through a Grant-in-Aid award. The authors would like to thank the occupants of the DC Nanogrid House for their patience during testing.