By Vince Wilson, Fangxing Li, Jesse Thornburg, Javad Mohammadi, and Justin Martinez

Energy Savings Through Refrigeration Load Control With Assessment of Commercial Potential

Lowering peak power and electricity bills through optimal demand scheduling.

THE WORLD IS CURRENTLY UNDERGOING AN energy transition driven by countries' efforts in decarbonizing their economies to both mitigate anthropogenic global warming and avoid costly fuel supply disruptions. Replacing fossil-fueled generators with renewable energy sources is a key method of decarbonization, but renewable power generation is

Digital Object Identifier 10.1109/MELE.2023.3348352
Date of current version: 29 February 2024

2325-5897/24©2024IEEE

intermittent and nonsynchronous. These factors decrease synchronous inertia and alter the availability of power on an electrical grid's supply side, threatening the grid's stability and reliability. Additionally, shortages in available power can lead to exorbitant electricity prices, threatening the finances of electricity consumers.

Challenges in Load Management

Demand response (DR) methods seek to address these effects by actively managing the level of electricity demand in response to the available power supply. DR can thus mitigate the ill effects of renewable energy sources on electrical grids and protect consumers from sudden price shocks.

One of the most prominent forms of DR is price-based DR (PBDR). PBDR programs enable loads to intelligently determine their operation in response to price signals sent by utilities or DR aggregators. The resulting DR action can be as simple as a load shutting off if electricity prices exceed a certain threshold to a load determining its daily operation schedule using an optimization function, usually to minimize the cost of electricity purchased. The latter is a very desirable form of PBDR, as it naturally shifts load operation to coincide with periods of low prices and high supply, leading to consistent savings for load operators and decreased demand when power is less available (thereby improving grid reliability). To maximize the effectiveness of load-shifting PBDR programs, the chosen loads must be highly deferrable and interruptible; i.e., the load must be able to delay its operation and (once started) cease operation on command. This limits suitable loads to those in which users will not notice interruptions in operation or those that have access to energy storage that enables them to continue operation without drawing power from the grid.

Given the previous limitations, commercial refrigeration systems have several advantages when used in DR. Because it takes time for refrigerated items to heat up, commercial refrigerators can be switched off for tens of minutes with no ill effects. The presence of variable-speed compressors, adjustable expansion valves, and active defrost units gives operators additional flexibility in planning a refrigerator's power consumption outside of a simple on–off switch. Typical commercial refrigeration systems possess a wide variety of sensors and network connectivity, which can aid automation efforts. Commercial refrigeration systems are common in terms of geographic distribution and constitute a considerable demand, consuming roughly 4% of the total electrical energy generated in the United States annually. Lastly, refrigeration is significantly more power intensive than air conditioning, giving PBDR program operators more controllable demand (kilowatts) per unit enrolled. All of these qualities combine to make commercial refrigeration an extremely attractive load type for use in DR.

DR programs have used the heating, ventilation, and air-conditioning systems of buildings to provide DR for many years now. Load shifting and peak demand reduction have been provided by preheating or cooling buildings at the request of grid operators. Fast frequency response and emergency generation capacity have been provided using automated remote-control systems installed in utility customers' air conditioners and water heaters. Being a related technology, refrigeration systems have been thoroughly investigated for DR but have yet to

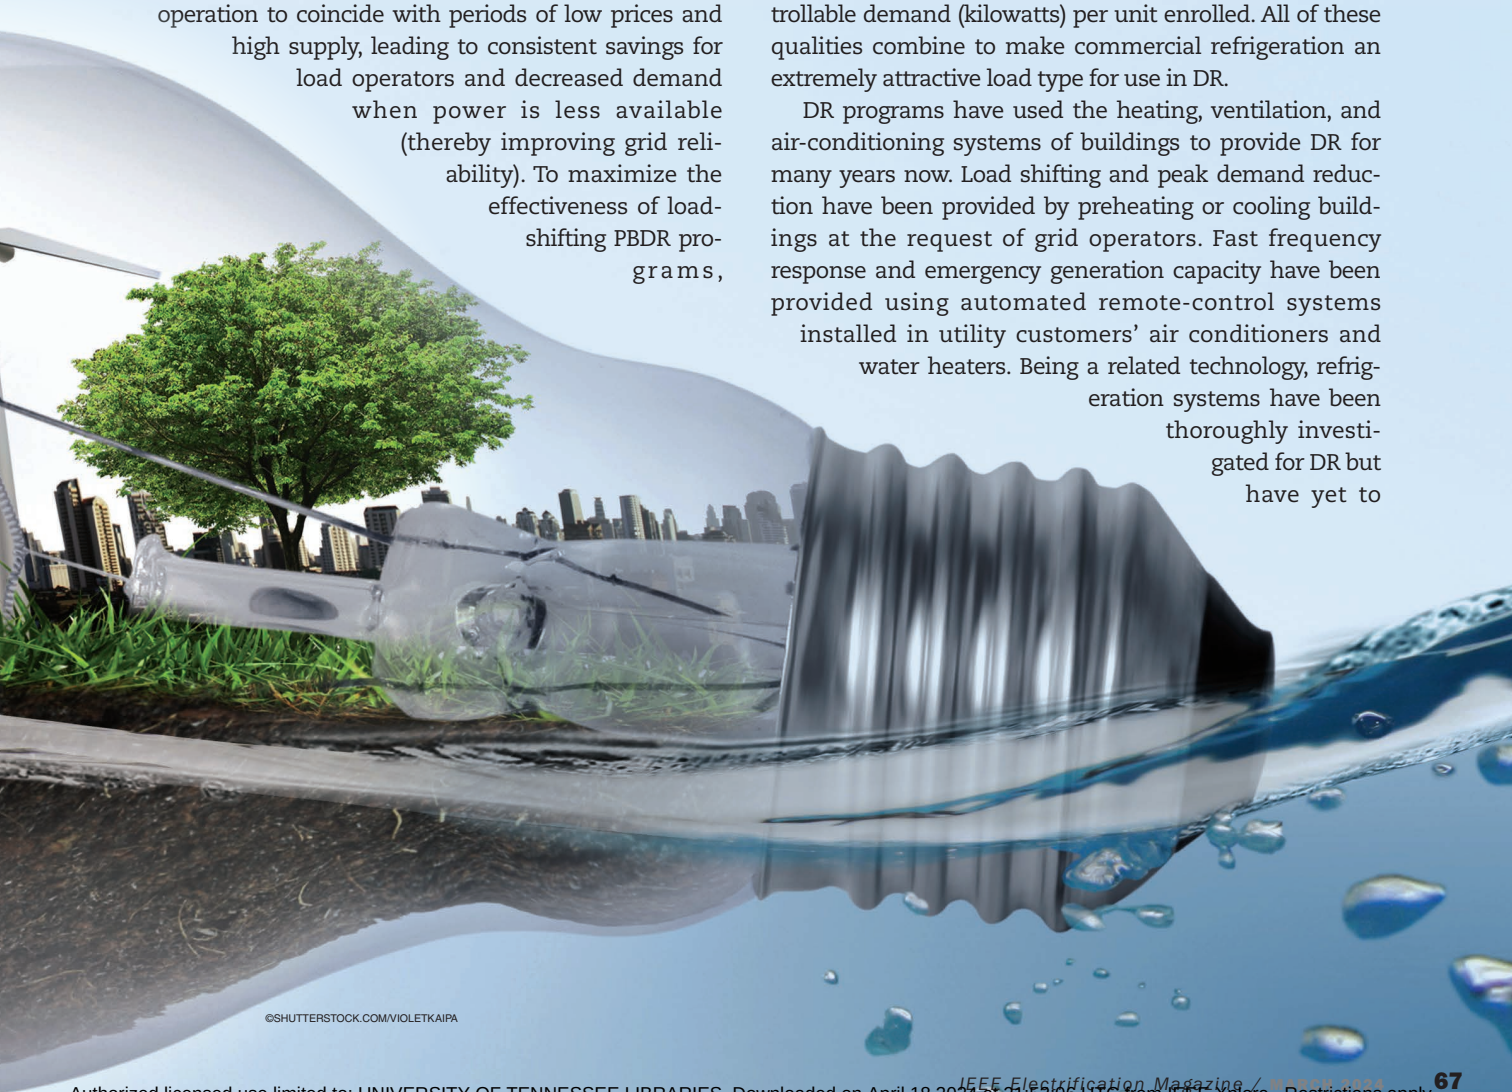

©SHUTTERSTOCK.COM/VIOLETKAIPA

see widespread enrollment. The contributions of this article are twofold. First, a simple thermal control prototype based on empirical measurements and suitable for deployment on low-computational-power devices is developed. Next, mixed-integer programming is used to examine the effects of daily variations in refrigerated stock on optimal compressor-control schedules and potential reductions in electricity costs. This mixed-integer program can be used to set an optimal daily schedule of compressor operation, which after a year of simulated operation was found to decrease electricity costs and energy consumption by 7.48% and 5%, respectively.

Testing Platform

To develop a realistic thermal model, a preexisting smart home and an automated demand response (ADR) demonstration project were repurposed into a testing platform. The platform was originally constructed to demonstrate the effectiveness of shifting refrigerator loads using the OpenADR communication standard as part of a joint project between the University of Tennessee and Grid Fruit LLC, funded by the Oak Ridge National Laboratory's RevV program with the goal of spurring technology innovation. The aim of the testing regime that followed was to collect enough data to determine the heating and cooling rates of a refrigerator when various levels of load mass are to be refrigerated. These data set was then used to construct a mathematical thermal model of the system for use in a scheduling mixed-integer linear program (MILP). Empirical data and curve fitting were used in the model construction for two reasons. First, the testbed refrigerator lacked documentation on several key values necessary for the use of differential heat-transfer equations. Second, it was desired that the final scheduling program be able to run on edge-computing devices with low computational power, such as a Raspberry Pi, preventing the use of ODE solvers.

Demonstration Platform Components and Topology

The testing platform was developed as a technology demonstrator for pairing the OpenADR demand response communication standard with commercial refrigerators and freezers. The platform consists of the following components: an Igloo FR3201 refrigerator, a central development PC, a Raspberry Pi 4, a UbiBot WS1 wireless thermometer, an IotaWatt Wi-Fi power monitor, and two Songle SRD-05VDC relays. Given that the refrigerator only consumes around 100 W when turned on, the relays would likely need to be uprated for full-scale use. Excluding the power monitor, which was only needed for testing, the total component cost is around US$133: US$35 for the Raspberry Pi, US$90 for the UbiBot thermometer, and US$8 for the relay boards. The development PC receives air temperature readings from the UbiBot thermometer and simulated locational marginal price (LMP) and ADR signals via the reading of XML files formatted in the OpenADR standard. The Raspberry Pi receives commands from the development PC and activates or deactivates the corresponding relays; the thermostat control relay enables or disables the stock thermostat while the compressor control relay directly turns the compressor on or off. Figure 1 shows the overall architecture of the system, with orange lines representing the flow of data and control signals and green lines showing the flow of electricity. Units meant for actual deployment would eliminate the development PC and receive temperature and ADR signals directly, as shown by the solid orange lines.

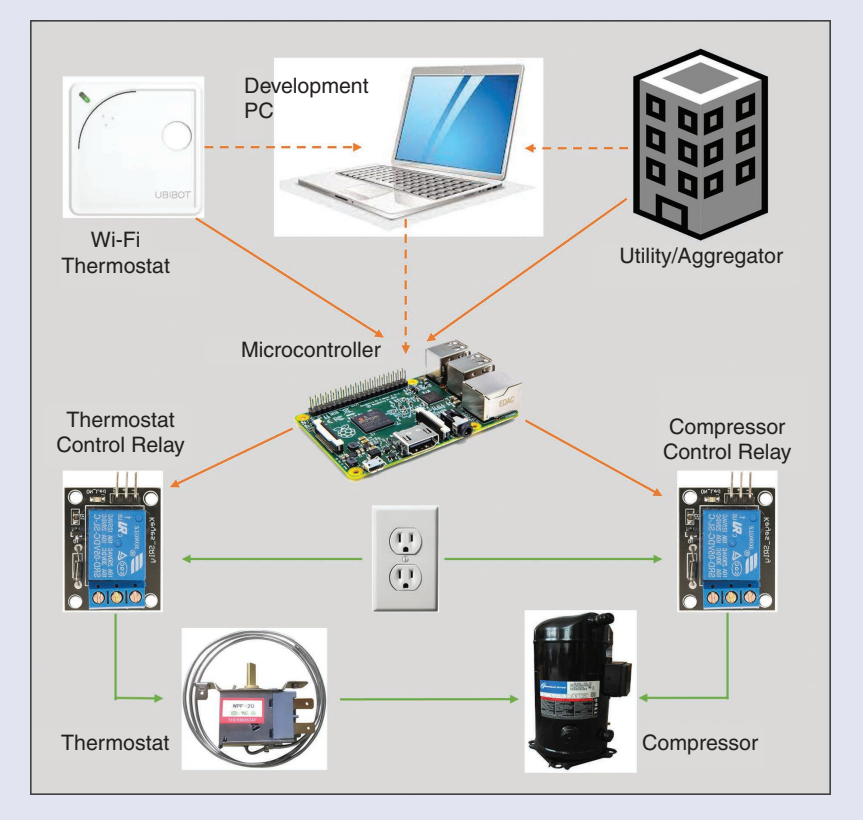

Figure 1. Test platform topology and information flows.

There are two control schemes that the system can use. The first involves simple dead-band logic: as long as the air temperature is between an upper and a lower bound, the compressor is turned off. As soon as the air temperature exceeds the upper bound, the compressor is turned on until the air temperature is below the

lower bound, at which point the compressor is turned off again. The second control scheme turns the compressor on or off at preset times, given that the current air temperature is between the upper and lower boundary temperatures. If the temperature is outside the boundaries, the compressor is turned on or off until the temperature is within bounds. The times when the compressor state is to be set can be determined by the user, an optimal scheduling MILP, or OpenADR signals.

Determining the System's Heating and Cooling Rates

To characterize the system's thermal behavior, heating and cooling curves were recorded with three different amounts of stock to be refrigerated. These amounts were 16 kg, 8 kg, and 0 kg of water, hereafter referred to as the full-, medium-, and no-stock conditions, respectively. To perform a cooling test, the requisite amount of room-temperature (21 °C) water for the given stock condition was placed into the refrigerator, and the refrigerator's compressor was set to a constant on state. The refrigerator's internal air temperature was then recorded every 15 s for the length of time it took for the internal air to cool to 0 °C from room temperature. Heating tests were conducted immediately following the conclusion of a cooling test by turning off the compressor and recording the internal air temperature until it returned to 21 °C. The U.S. Food and Drug Administration defines refrigeration temperature as being less than or equal to 5 °C; as such, the collected data were truncated to include values only within that temperature range. Linear curve fits were performed on the remaining data to obtain heating and cooling rates for the given stock conditions in the 5 °C to 0 °C temperature range. See Table 1.

Thermostatic and Optimized Scheduling Simulation Methods

Using the data collected via the experiments discussed in the previous section, two different models were constructed that simulate the system's temperature and power consumption over the course of 24 h. Datasets containing historical outside air temperature, real-time (RT) LMP, and day-ahead (DA) LMP were obtained, and simulations were repeated with varying daily datasets to model a year of continuous system operation. The first set of simulations provides a baseline where the refrigerator is only capable of hysteretic control from the factory-installed thermostat. The second set uses an optimal scheduling MILP to directly control the compressor and minimize electricity costs for the current day. Both simulation types share several elements, such as variations in refrigeration stock throughout the day and outside temperatures affecting system power consumption.

Elements Present in Both Simulations

There are several commonalities between either simulation type, the first of which is the LMP and outdoor temperature datasets. Hourly DA and RT LMP values from 1 January 2022 to 1 January 2023 were obtained using the

TABLE 1. Curve-fit equations and corresponding heating or cooling rates.

Test Condition	Curve-Fit Equation	$T\Delta$, °C/s	r^2 Value
Full-stock cooling	–0.00009*t + 4.3496	–0.00009	0.9443
Medium-stock cooling	–0.0001*t + 3.9998	–0.0001	0.9218
No-stock cooling	–0.0029*t + 4.9556	–0.0029	0.9911
Full-stock heating	0.0002*t + 0.4274	0.0002	0.9886
Medium-stock heating	0.0004*t + 1.1054	0.0004	0.8803
No-stock heating	0.0016*t + 0.1949	0.0016	0.9622

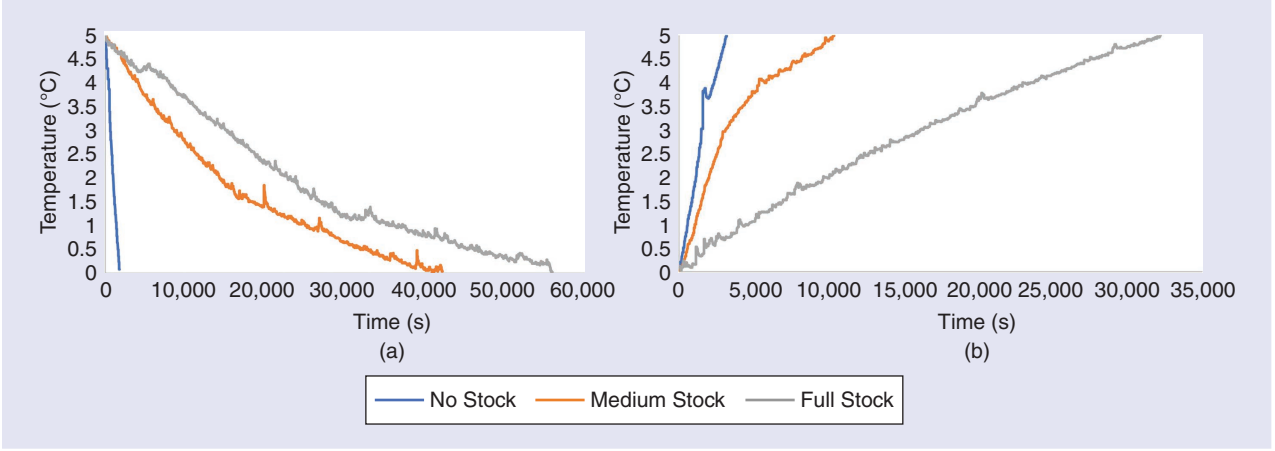

Figure 2. Temperature-over-time data for full-, medium-, and no-stock conditions. (a) Cooling test results. (b) Heating test results.

$$T(i+1) = T\Delta * t + T(i) \quad (1)$$

$$T = T_1 + (X - X_1) * \frac{(T_2 - T_1)}{(X_2 - X_1)} \quad (2)$$

$$\text{COP} = \frac{Q}{W} = K * \eta \quad (3)$$

$$K = \frac{T_C}{T_H - T_C} \quad (4)$$

$$P = \frac{Q_{\text{base}}}{\eta * \frac{T_C}{T_H - T_C} * 3{,}600} \quad (5)$$

T	System Interior Air Temperature
$T\Delta$	Heating or Cooling Rate
t	Length of Timestep
X	Mass of Refrigeration Stock
COP	Refrigerator Coefficient of Performance
Q	Heat Removed From System
W	Work Put into System
K	Carnot Efficiency
η	Refrigerator Efficiency (Assumed Value)
T_C	Evaporator Surface Temperature
T_H	Outdoor Air Temperature
Q_{base}	Refrigerator Steady State Cooling Power (Calculated)
P	Power Required for Cooling

Figure 3. Equations used in both simulation types.

PJM Data Miner tool for Athenia, NJ, USA. Hourly outdoor air temperature readings over the same time period were obtained from Newark Liberty International Airport using the Iowa Environmental Mesonet (IEM) service. Athenia was chosen as the simulated location because of its distribution voltage level (26 kV) and urban location with significant restaurants and grocers, making for a good representation of commercial refrigeration's typical operating environment. Lastly, each daily simulation is made up of 144 time steps or 10 min/time step, starting at 12 a.m. on the current day and ending at the same time the following day. The length of time steps is chosen to ensure that the compressor remains in each state for a long enough period to avoid short cycling; turning a compressor on and off in rapid succession consumes excess power and can drastically shorten the unit's life span.

As can be seen in Table 1, the heating and cooling rates are mostly linear, with an average r^2 value of 0.95 (r^2 is a statistical measure showing how well the data fits a regression model). Figure 3 presents the equations used in both simulation types. Given this, (1) in Figure 3 is used in each simulation to approximate future internal air temperatures by multiplying the appropriate heating or cooling rate by a certain amount of time and adding that value to the current internal air temperature. To simulate future temperatures for refrigeration stock levels that do not have recorded curves, linear interpolation, (2), is used. Here, X is the mass of the refrigerated stock for the off-curve condition and T is the predicted temperature. X_1, T_1, X_2, and T_2 are the refrigerated stock mass and temperature from the corresponding curve. For example, if X equals 4 kg, X_1 and T_1 would be from the no-stock condition and X_2 and T_2 from the medium-stock condition. The coefficient of performance (*COP*) (3) is the efficiency of a reverse heat engine. The refrigerator's *COP* is affected by the temperature differential between the outside air and the condenser. If the refrigerator's evaporator temperature is held constant, higher outside air temperatures decrease *COP* and vice versa, causing significant seasonal variations in system power consumption. *COP* can also be thought of as the product of the Carnot efficiency, K and η; K gives the maximum *COP* possible for a reverse heat engine, while η is an assumed percentage value for the system's actual efficiency due to its physical properties.

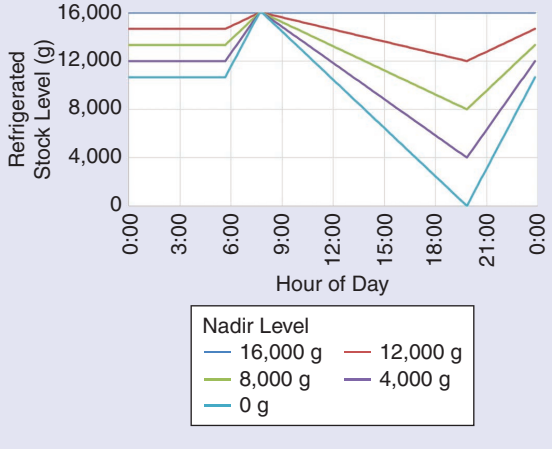

Figure 4. Refrigerated stock variation patterns.

To find the change in *COP*, we must first determine the baselines of several values. COP_{base} is set at 1.625 to match average *COP* values for commercial refrigerators. T_H is 285.59 K since condensers are typically rooftop mounted, and that was the average outdoor temperature in the contiguous United States for 2020. T_C is 269.75 K; this was found by measuring the evaporator surface temperature with an infrared thermometer. K_{base} and η are then calculated to be 17.4 and 0.07, respectively. The compressor in the test system draws 38 W of power at steady state; using this value, $Q_{\text{base}} = COP_{\text{base}} * 38 * 3600 = 171{,}000$ J of heat are removed over the course of an hour. It should be noted that the values of η_{base} and Q_{base} have not been empirically confirmed nor do they need to be since we are only concerned with the differences between *COP* values at a certain T_H and not the absolute *COP* values. To maintain a constant rate of cooling, the required power must change

as T_H does. This power can be calculated using (5) in Figure 3, with η and T_C equaling 0.07 and 269.75 K, respectively, and T_H being the outdoor air temperature at the desired time.

The amount of refrigeration stock throughout the day is either held constant or varied according to a set pattern depending on the test scenario. This pattern is as follows: at 8 a.m., a grocery store has just opened, and the refrigerators are assumed to be fully loaded with 16 kg of stock. The amount of refrigerated stock then decreases linearly over the course of the day, reaching a nadir at 8 p.m. From 8 p.m. to 12 a.m. and from 6 a.m. to 8 a.m., the refrigerators are restocked at a constant rate until the stock levels reach 16 kg. The stock levels over time for each test scenario can be seen in Figure 4. By implementing this stock variation pattern, the loss of thermal inertia that would occur as customers remove products from the refrigerators over the course of a day can be accounted for.

In our case, the objective function seeks to minimize the cost of electricity purchased each day while keeping the refrigerator's internal air temperature within safe boundaries.

Thermostatic Baseline Simulation

To determine the effect of load variation on optimal scheduling, a baseline for comparison was established by simulating the operation of the test fridge under the control of a basic thermostat using a rule-based model. In this case, the refrigerator is following the simple dead-band logic described in the section "Demonstration Platform Components and Topology" and can easily be simulated using loops and conditional statements. See Algorithm 1. First, T_{heat} and T_{cool}, which are the possible internal air temperatures in the next time step if the system heats or cools, respectively, are calculated using the current refrigerated stock levels and (1) and (2) of Figure 3. Next, if T_{heat} is greater than 5 °C, it means that the dead band has been broken, and the system must start cooling. The compressor state, p, is set to 1, T(i + 1) is set equal to T_{cool}, and P is calculated using (5). Similarly, if T_{cool} is less than 1 °C, p and P both equal 0 and T(i + 1) is set equal to T_{heat}. The lower dead-band boundary is set to 1 °C to ensure that the water remains liquid and the additional complexity introduced by the phase change from liquid to solid can be avoided.

If the current temperature is within the dead band, the system could either be in a cooling cycle or heating up after a cooling cycle has finished. If the former, that means p(i-1) would be 1 and if the latter, p(i-0) would be 0. When either of these cases is true, that means p(i) must be set to p(i-1) and P and T(i + 1) assigned the corresponding values according to the pattern from the first two conditions. The only exceptions to these rules are the simulation's initial conditions. In this case, the first day has T(0) set to 2 °C and p and P set to 0. For the following days, the initial conditions are set equal to the previous day's final states. It can be seen in Figure 5 that cooling cycles become more frequent later in the day as refrigerated stock reaches its nadir and the system has the lowest amount of thermal inertia.

ALGORITHM 1: Thermostatic simulation pseudocode.

For i in 0 through 144:
  Calculate T_{heat} and T_{cool} using (1) and (2).
  If $T_{heat} \geq 5$ °C:
    $T(i + 1) = T_{cool}$, p(i)=1, P(i) calculated using (5).
  ElseIf $T_{cool} \leq 1$ °C:
    $T(i + 1) = T_{heat}$, p(i) and P(i)=0
  ElseIF p(i-1) =1:
    $T(i + 1) = T_{cool}$, p(i)=1, P(i) calculated using (5).
  Else:
    $T(i + 1) = T_{heat}$, p(i) and P(i)=0

Optimal Scheduling Program Simulation

Optimal device scheduling in response to price signals, weather conditions, or other factors is a well-explored topic. In our case, the objective function seeks to minimize the cost of electricity purchased each day while keeping the refrigerator's internal air temperature within safe boundaries. As such, only a single decision variable, p, which describes the on–off state of the compressor is needed. Each daily optimization was solved

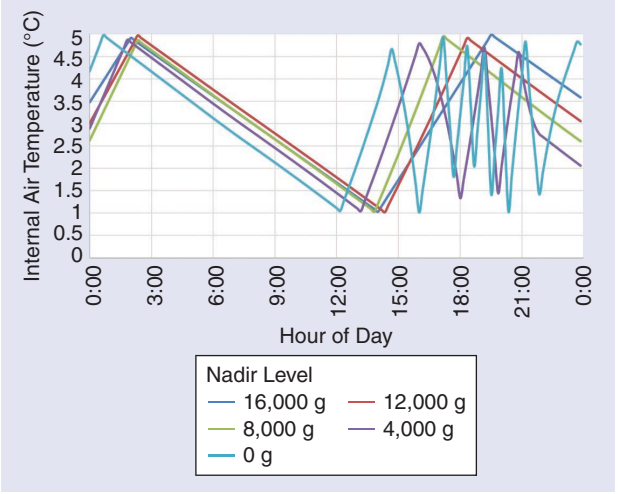

Figure 5. Sample of temperature profiles from thermostatic simulations for 2 July 2022.

using Gurobi Optimizer running in a Python 3.9.12 programming environment. Figure 6 presents the optimal scheduling MILP. Figure 7 describes the usage of heating and cooling data in the thermostatic and MILP simulations. As shown in Figure 8, the schedule optimizing MILP produces schedules that account for hourly variations in price by precooling the refrigerator during low-price, low-compressor-power periods so that the compressor can be disabled during high-price, high-compressor-power periods. Mornings and nights are the typical low-price periods with afternoons and evenings having higher prices due to greater demand. This trend is generally the same for compressor power as mornings and nights usually have lower outdoor temperatures compared to midday, barring heat waves or cold snaps.

The result is a nonlinear relationship where the greater ease of cooling conflicts with increased heating rates to determine how much energy is needed to maintain a certain temperature.

Simulation Results

A total of 20 scenarios were simulated to examine the influence of stock variation and how it may affect optimal scheduling algorithms. These simulation runs were divided into four groups of five. The first two groups examined the performance of the thermostatic and optimal scheduling control methods with the amount of refrigerated stock held constant. The stock was held constant in each test but decreased in increments of 4 kg between tests from 16 kg to 0 kg. The third and fourth groups also examined the difference between thermostatic and optimal scheduling control but with refrigeration stock now varying throughout the day. For each test, the amount of refrigerated stock followed one of the patterns shown in Figure 4. Values of interest are the cost of electricity based on the DA LMP, the amount of energy used during a given period, and peak power consumption.

$$\text{Minimize: } \Sigma_1^{144}\, p(i)*\text{LMP}_{\text{DA}}(i)*P(i)*\frac{600}{3{,}600}*10^{-5} \quad (6)$$

Subject to:

$$T(i) \geq 1\ ^\circ\text{C},\ i \in [1{,}144] \quad (6a)$$

$$T(i) \geq 5\ ^\circ\text{C},\ i \in [1{,}144] \quad (6b)$$

$$p(i) == 0 >> T_{\text{med}}(i) == T(i) + 0.0004 * 600,\ i \in [1{,}144] \quad (6c)$$

$$p(i) == 1 >> T_{\text{med}}(i) == T(i) - 0.0001 * 600,\ i \in [1{,}144] \quad (6d)$$

$$p(i) == 0 >> T_{\text{alt}}(i) == T(i) + 0.0002 * 600,\ i \in [\text{stock}(i) \geq 8{,}000] \quad (6e)$$

$$p(i) == 1 >> T_{\text{alt}}(i) == T(i) - 0.00009 * 600,\ i \in [\text{stock}(i) \geq 8{,}000] \quad (6f)$$

$$T(i{+}1) == T_{\text{med}}(i) + (\text{stock}(i) - 8{,}000) * \frac{T_{\text{alt}}(i) - T_{\text{med}}(i)}{8{,}000},\ i \in [\text{stock}(i) \geq 8{,}000] \quad (6g)$$

$$p(i) == 0 >> T_{\text{alt}}(i) == T(i) + 0.0016 * 600,\ i \in [\text{stock}(i) < 8{,}000] \quad (6h)$$

$$p(i) == 1 >> T_{\text{alt}}(i) == T(i) - 0.0029 * 600,\ i \in [\text{stock}(i) < 8{,}000] \quad (6i)$$

$$T(i{+}1) == T_{\text{alt}}(i) + (\text{stock}(i) * \frac{T_{\text{med}}(i) - T_{\text{alt}}(i)}{8{,}000},\ i \in [\text{stock}(i) \geq 8{,}000] \quad (6j)$$

Decision Variables	
p	Compressor State

Parameters	
LMP_{DA}	Day-Ahead Locational Marginal Price
p	Compressor Power Found Using (5)
T	System Air Temperature for Current Stock
T_{med}	Medium-Stock Air Temperature
T_{alt}	Air Temperature for Full- or No-Stock Conditions
stock	Mass of Refrigeration Stock

Figure 6. Optimal scheduling MILP.

Analysis of One Year of Simulated Operation Under Both Methods

To determine the potential of commercial refrigeration systems as DR resources, two factors are examined. First, the effect of constant and variable refrigeration stock is analyzed to draw some general conclusions on how stock levels impact system operation. Next, the performance of thermostatic and optimal scheduling simulations is compared in terms of electricity cost, energy usage, and peak power demands (see Figure 9). The goal is to determine if the optimal scheduling program can produce consistent reductions in electricity cost when the thermal performance is derived experimentally from a real-world test system. In addition, energy usage and peak power demand are examined to determine if optimal scheduling programs may have deleterious effects on the refrigeration sector's overall energy consumption or the grid's ability to meet demand during peak demand periods.

Influence of Refrigerated Stock Variation

In both thermostatic and optimal scheduling, for constant-load scenarios in the range of 16 kg to 8 kg, the cost, energy usage, and peak power increase as the amount of refrigerated stock decreases; the yearly DA cost, yearly energy usage, and average peak power increased on average 13%, 9%, and 1.3%, respectively. In the range of 8 kg to 0 g, the trend reverses with a significant drop in prices and energy usage going from 8 kg to 4 kg and a smaller decrease going from 4 to 0 kg. The amount of energy used by the system is determined by both the length and frequency of cooling cycles, which are themselves determined by the system's thermal inertia. High-inertia cooling cycles are long but infrequent, while low-inertia cooling cycles are the opposite. In other words, the lower the thermal inertia, the less energy is needed to cool, and so cooling cycles become shorter. On the other hand, less energy is needed to heat the system as well, and so more cooling cycles are needed to maintain a certain temperature. The result is a nonlinear relationship where the greater ease of cooling conflicts with increased heating rates to determine how much energy is needed to maintain a certain temperature.

A similar result can be observed when it comes to simulations using the refrigerated stock variation patterns, though the magnitude of increases in the 8-to-16-kg range is smaller: 9%, 6%, and 0.55% for yearly DA cost, yearly energy usage, and average peak power, respectively. This can be explained by the mass of refrigerated stock being greater on average in variable-stock scenarios than in constant-stock scenarios. For example, the 8-kg nadir scenario has an average refrigerated stock mass throughout the day of 12.351 kg versus a constant 8 kg for the constant-stock scenario. This can be taken to the logical extreme with the 0-kg nadir scenario, which has an average stock mass of 8.7 kg versus 0 kg for the constant stock. Because the variable-stock scenarios must always be fully restocked to 16 kg, they will always have more thermal inertia to work with and will more closely resemble higher constant-stock scenarios. This is relevant to real-world operation as the amount of stock at a given time can drastically alter a system's heating and cooling requirements, as shown in Figure 10. A practical recommendation then would be to sync up the restock periods to maximize the system inertia before the high-price periods.

A practical recommendation then would be to sync up the restock periods to maximize the system inertia before the high-price periods.

Energy and Financial Savings

Possible reductions in cost, energy usage, and peak power can be determined by comparing these values between the relevant thermostatic and optimal scheduling simulations. In general, optimal scheduling of commercial refrigerators leads to reductions in electricity cost and energy usage and a small increase in peak power usage. For constant-stock simulations which are idealistic, optimal scheduling can reduce the yearly DA cost and yearly energy usage by 8.6% and 5.36%, while the peak power consumption barely increases by 0.14%. For variable-stock simulations, the 16-kg scenarios are excluded since, by nature of how the variation patterns are created, they actually maintain a constant stock of 16 kg. For the remaining four simulation types of

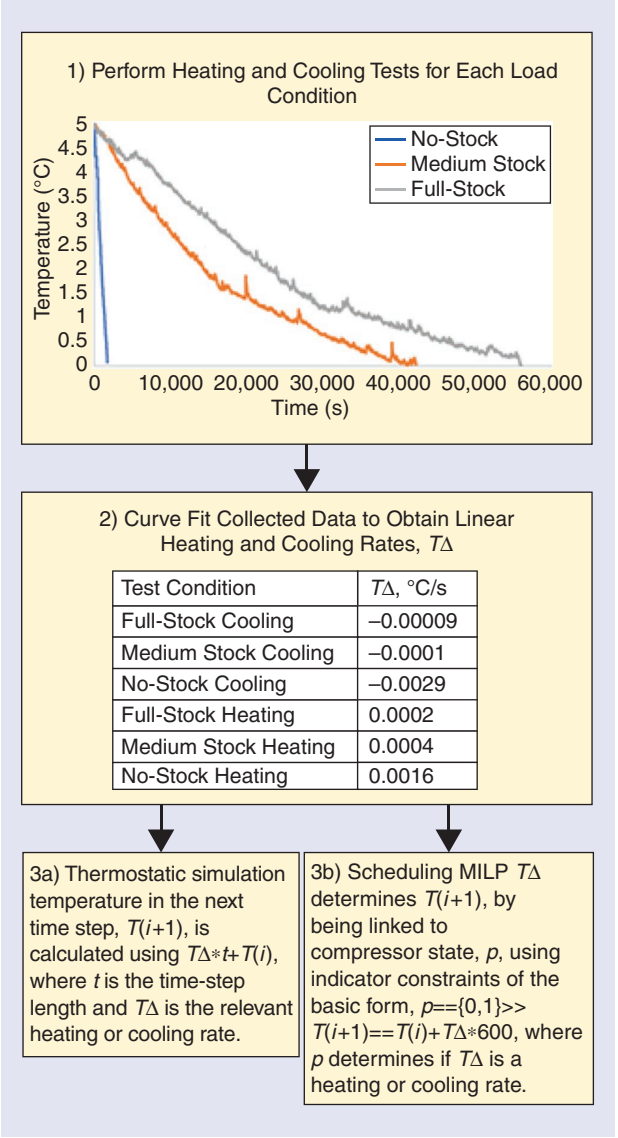

Figure 7. Usage of heating and cooling data in thermostatic and MILP simulations.

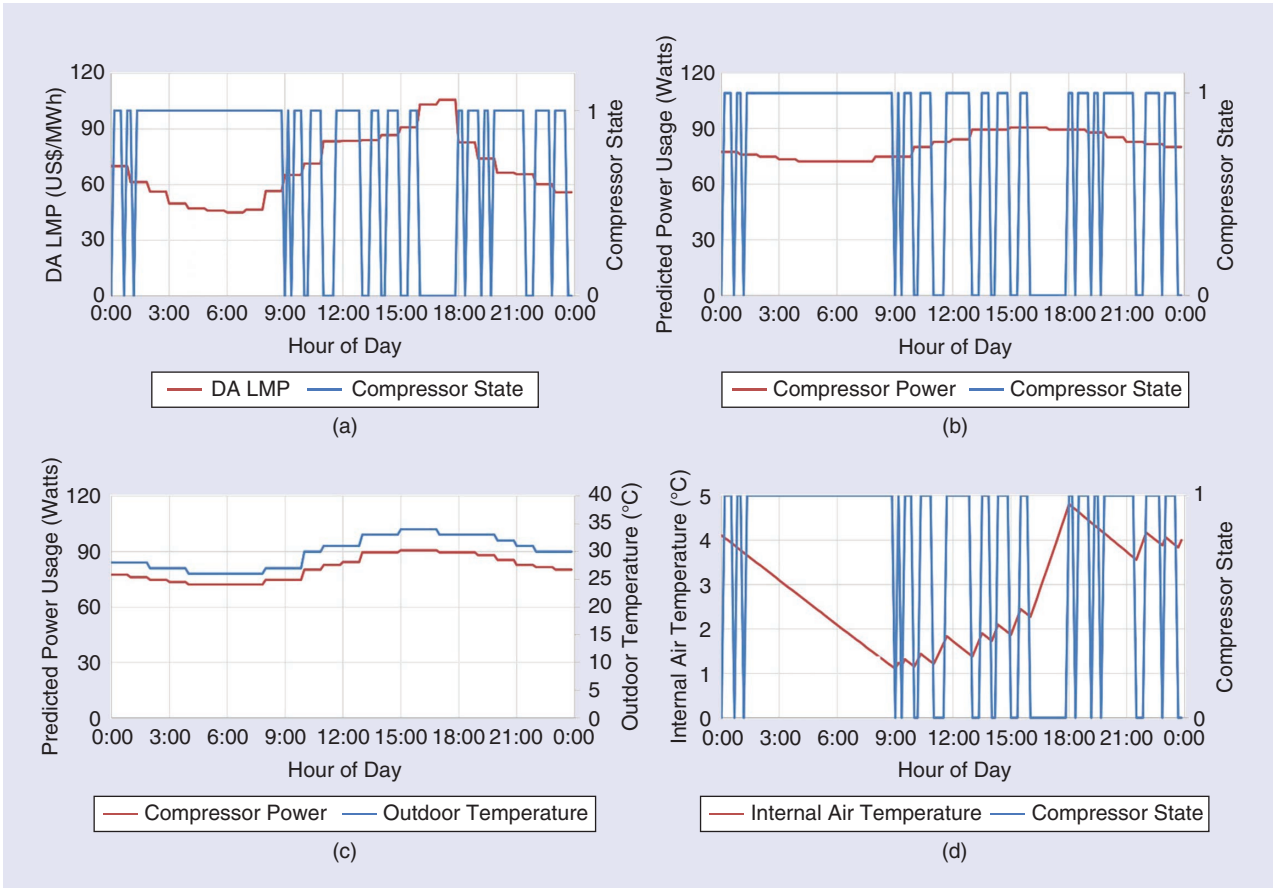

Figure 8. Optimal schedule parameters versus driving influences for 2 July 2022. (a) Compressor state versus price. (b) Compressor state versus predicted power usage. (c) Outdoor temperature versus predicted power usage. (d) Internal air temperature versus compressor state.

realistic variable stocks, the optimal scheduling algorithm reduces the average yearly DA cost by 7.48%, decreases the average yearly energy consumption by 5.00%, and increases the average peak power consumption by 0.72%. It should be noted that most of the reductions in energy consumption occur in the 4- and 0-kg nadir scenarios as the change in energy consumption for the 12- and 8-kg nadir scenarios is marginal.

The reductions in electricity costs occur as the optimal scheduling algorithm preemptively cools the system to a low temperature before a high price period so the system can heat up over said high price period and not require the compressor to be on. The reduction in energy consumption is the result of the optimal scheduling algorithm being able to carry out variable-length cooling cycles; when operating under thermostatic control, the system will always cool from high to low temperature extremes and then allow itself to heat back up to 5 °C or near it. It may be more energy efficient to maintain the temperature at a certain level instead of allowing this heating to occur, which can be done using short and rapid cooling cycles, as seen in Figure 10. Lastly, the optimal scheduling algorithm increases peak power consumption since daily outdoor temperatures usually peak slightly before electricity prices. This means that the preemptive cooling usually takes place during the hottest part of the day when the condenser is the least efficient.

Commercial refrigeration systems typically take the form of refrigerated cabinets and walk-in units, of which there were 706,000 and 837,000 buildings equipped with each, respectively, as recorded in the 2018 U.S. Energy Information Administration (EIA) Commercial Buildings Energy Consumption Survey; cabinets used 499 billion kWh of electricity while walk-in units used 546 billion kWh, costing their operators US$47.73 billion and US$51.64 billion annually to operate. Assuming equal levels of savings are possible with full-sized commercial systems as the simulated testing platform for both cabinets and walk-in units, each building with cabinets could expect to save US$5,056.94 (= $47.73 * 10^9 * 0.0748/706{,}000$) and reduce their energy consumption by 35.34 MWh (= $499 * 10^9 * 0.05/706{,}000$ kWh). Using similar calculations, buildings with walk-in units can expect US$4,614.90 and 32.62 MWh in cost and energy usage reductions. Reductions in electricity consumption would also lower a building's carbon emissions; the production of a kilowatt hour of electricity releases 0.387 kg of CO_2 per EIA estimates. Optimal scheduling systems could therefore stop 13.68 kilotons and 12.62 kilotons of CO_2, respectively, from entering the atmosphere per building

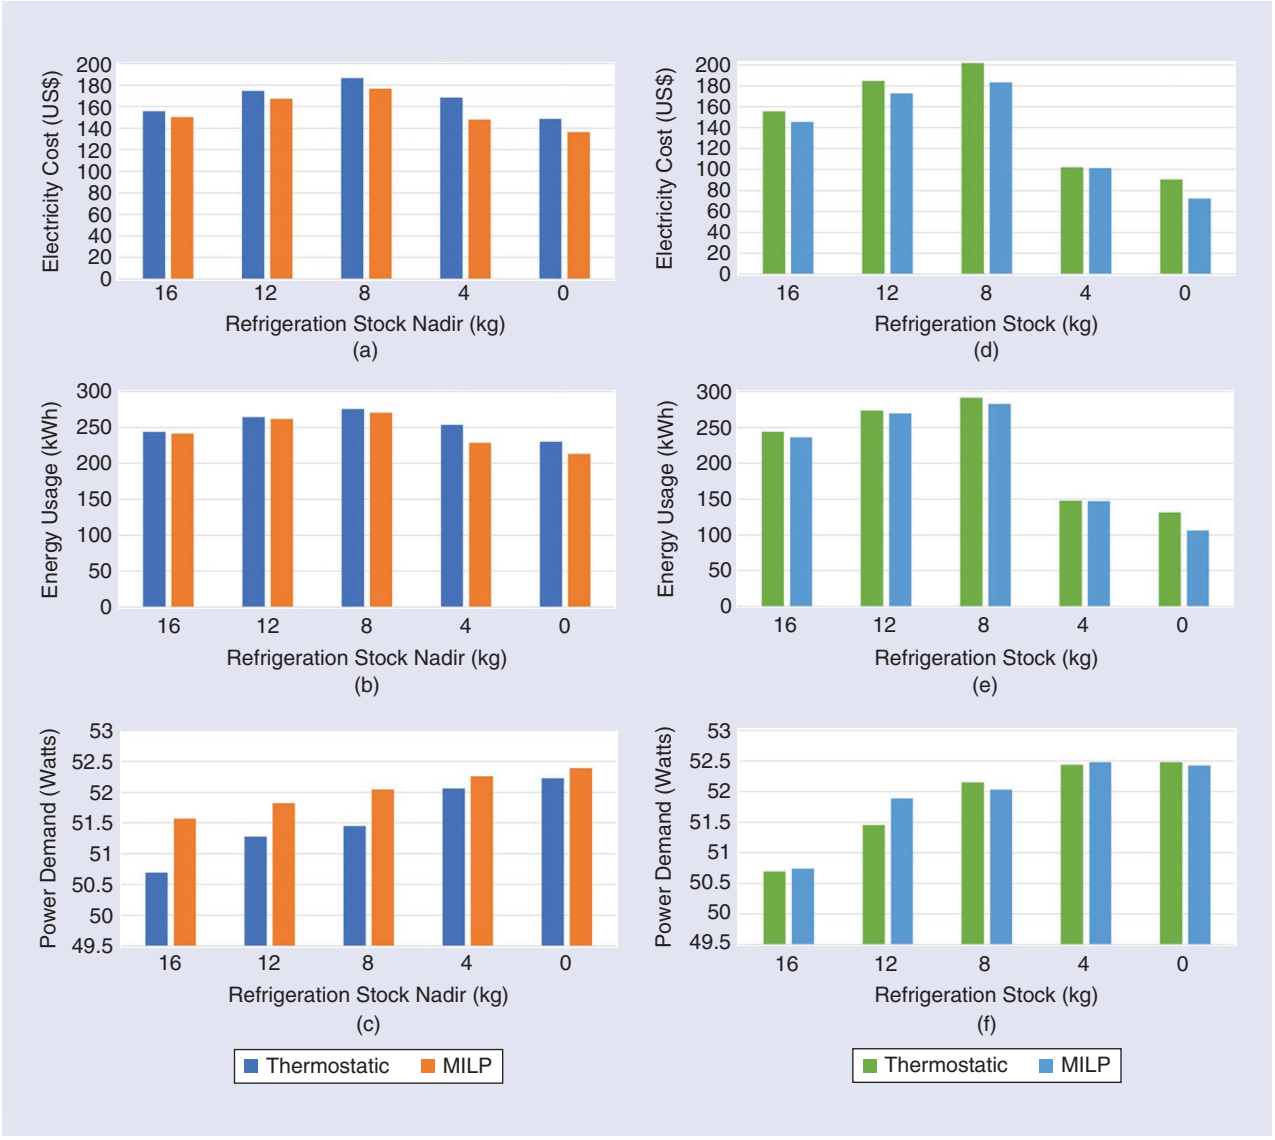

Figure 9. Charts displaying various results of simulated scenarios. (a) Yearly electricity cost for variable-stock simulations. (b) Yearly energy usage for variable-stock simulations. (c) Average daily peak demand for variable-stock simulations. (d) Yearly electricity cost for constant-stock simulations. (e) Yearly energy usage for constant-stock simulations. (f) Average daily peak demand for constant-stock simulations.

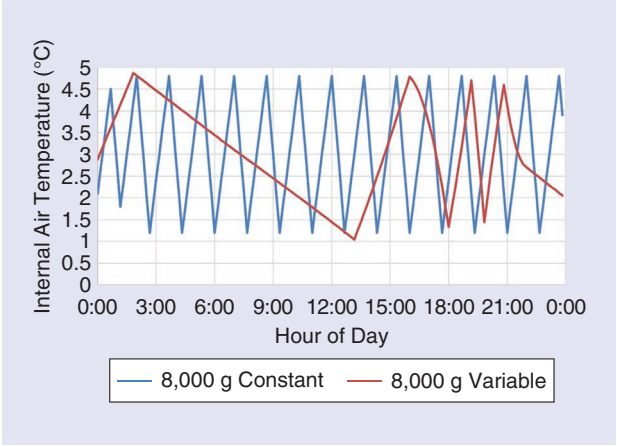

Figure 10. Thermostatic temperature profiles for constant- and variable-stock scenarios on 2 July 2022.

due to the operation of cabinets and walk-in units. Assuming an average commercial building has 10 of either unit type, the total retrofit cost would be US$1,330; this would provide a return of 2.8 [= (5,056.94–1,330)/1,330] and 2.47 [= (4,614.9–1,330)/1,330] for cabinets and walk-in units respectively. If implemented on a nationwide basis, the use of optimal load scheduling for buildings with commercial-grade refrigeration cabinets and walk-in units could result in a net financial savings of US$7.43 billion, 52.25 million MWh lower electricity demand, and 20,220 megatons of CO_2 emission reductions.

Caveats to the Savings Analysis

There are several factors associated with deployment on full-sized commercial systems that suggest the prior

savings analysis is underestimated. First, the current analysis was done using an optimal scheduling MILP that can only control a constant-speed compressor since that is the only power-consuming component in the current test platform. In reality, full-sized commercial systems commonly have additional power-consuming components, such as defrosters, variable-speed compressors, and interior fans. By factoring in the presence of these features, future optimal scheduling MILPs would have significantly more control options for how and when power is being used. Second, full-sized commercial systems will have a much greater thermal inertia because their larger internal volume allows for more refrigeration stock to be held inside. While the available surface area for heat transfer on the outside of the refrigerator is also larger in full-sized systems, increases in volume always outpace increases in surface area. This ensures that inertia gains from scaling up a system will be greater than the increased exterior heating load. The additional thermal inertia increases the system's flexibility, allowing for deeper precooling and longer periods where the compressor can be off. These two factors mean it is likely that the potential cost savings could be boosted into the 10%–30% range predicted in the existing literature.

Conclusions

An analysis of the potential reductions in energy consumption and electricity cost from implementing an optimal scheduling algorithm in commercial refrigeration systems was estimated using empirical heating and cooling data and simulated systems. Heating and cooling data of a real-world test refrigeration system were collected and used to construct two models of the refrigeration system; one system operated under thermostatic controls, and the other was intelligently scheduled using linear optimization to minimize electricity costs in response to DA hourly LMPs. A year of operation was simulated for both systems with varying levels of load in response to historical weather and price data. The optimal scheduling algorithm reduced electricity cost and energy consumption by 7.48% and 5%, respectively, while slightly increasing the peak power demanded by the system. These results are likely underestimated since the scheduling algorithm is not able to control other components, such as defrosters and variable-speed compressors, necessitating further analysis to be more certain of the savings potential of commercial refrigeration.

Acknowledgment

For this research, Grid Fruit LLC provided the fundamental refrigeration management concept, while researchers conducted the algorithm development and testing. This material is supported by a RevV grant from Tennessee and the Oak Ridge National Laboratory, and Grid Fruit's contribution is based upon work supported by the U.S. Department of Energy, Office of Science, and Office of Building Technologies Office, under Award Number DE-SC0020822.

For Further Reading

Q. Hu and F. Li, "Hardware design of smart home energy management system with dynamic price response," *IEEE Trans. Smart Grid*, vol. 4, no. 4, pp. 1878–1887, Dec. 2013, doi: 10.1109/TSG.2013.2258181.

Y. Du, F. Li, K. Kurte, J. Munk, and H. Zandi, "Demonstration of intelligent HVAC load management with deep reinforcement learning: real-world experience of machine learning in demand control," *IEEE Power Energy Mag.*, vol. 20, no. 3, pp. 42–53, May/Jun. 2022, doi: 10.1109/MPE.2022.3150825.

C. Goodman, J. Thornburg, S. K. Ramaswami, and J. Mohammadi, "Load forecasting of food retail buildings with deep learning," in *Proc. IEEE PES Innov. Smart Grid Technol. Conf. - Latin America (ISGT Latin America)*, Lima, Peru, 2021, pp. 1–5, doi: 10.1109/ISGTLatinAmerica52371.2021.9543085.

J. Mohammadi and J. Thornburg, "Connecting distributed pockets of energy flexibility through federated computations: limitations and possibilities," in *Proc. 19th IEEE Int. Conf. Mach. Learn. Appl. (ICMLA)*, Miami, FL, USA, 2020, pp. 1161–1166, doi: 10.1109/ICMLA51294.2020.00186.

S. M. N. Hasnaeen, S. Rahman, and M. F. Uddin, "Load scheduling of a refrigerated warehouse with homogeneous compressors under dynamic pricing," in *Proc. 1st Int. Conf. Adv. Sci., Eng. Robot. Technol. (ICASERT)*, Dhaka, Bangladesh, 2019, pp. 1–6, doi: 10.1109/ICASERT.2019.8934558.

M. Glavan, D. Gradišar, I. Humar, and D. Vrančić, "Refrigeration control algorithm for managing supermarket's overall peak power demand," *IEEE Trans. Control Syst. Technol.*, vol. 27, no. 5, pp. 2279–2286, Sep. 2019, doi: 10.1109/TCST.2018.2853739.

M. Mohammadi, J. Thornburg, and J. Mohammadi, "Towards an energy future with ubiquitous electric vehicles: barriers and opportunities," *Energies*, vol. 16, no. 17, p. 6379, 2023, doi: 10.3390/en16176379.

M. Mohammadi and J. Thornburg, "Strategic EV charging and renewable integration in Texas," in *2024 IEEE Texas Power and Energy Conf. (TPEC 2024)*, College Station, TX, USA, 2024.

Biographies

Vince Wilson (vwilso17@utk.edu) is a Ph.D. student at the University of Tennessee, Knoxville, TN 37996 USA.

Fangxing (Fran) Li (fli6@utk.edu) is a John W. Fisher Professor at the University of Tennessee, Knoxville, TN 37996 USA.

Jesse Thornburg (jesse@gridfruit.com) is an assistant teaching professor at Carnegie Mellon University Africa and the CEO of Grid Fruit, LLC, Austin, TX 78758 USA.

Javad Mohammadi (javadm@utexas.edu) is an assistant professor at the University of Texas, Austin, TX **78712** USA.

Justin Martinez (BitterSteel152@gmail.com) is with the Cubic Corporation, Knoxville, TN 37932 USA.