Demand response capabilities of industrial refrigerated warehouses: Experiences in practical implementation

Andrea Mammoli
Electric Power Research Institute
3420 Hillview Ave
Palo Alto, CA 95304
USA
amammoli@epri.com

Don Shirey
Electric Power Research Institute,
942 Corridor Park Blvd
Knoxville, TN 37932
USA
dshirey@epri.com

Colin Lee
Electric Power Research Institute
3420 Hillview Ave
Palo Alto, CA 95304
USA
colee@epri.com

Ammi Amarnath
Electric Power Research Institute
3420 Hillview Ave
Palo Alto, CA 95304
USA
aamarnath@epri.com

Keywords

buildings, demand response, industrial energy saving, industrial internet of things (IIoT)

Abstract

As the amount of renewable electric power generation increases, so does the need to find end uses that can vary their demand to follow supply availability. Storage capacity inherent in certain end uses can be leveraged to enable flexible demand response (DR) to offset power generation intermittency. One such opportunity is in industrial refrigerated warehouses, which are used by logistics operations for storing food products as they move from production facilities to consumers. California hosts the largest amount of refrigerated space in the U.S., with refrigeration load of about 120,000 tons of refrigeration (TR) [422 MW_{th}] and associated electric power consumption of about 360 MW. A recent study by Lawrence Berkely National Laboratory estimates the DR potential of refrigerated warehouses at 0.575 kW / TR [0.163 kW_e/kW_{th}] by storing thermal energy in the refrigerated/frozen products, corresponding to at least 69 MW_e. An EPRI-led team investigated the potential of automated DR at a large industrial refrigerated warehouse in California. Using open standards based on the OpenADR 2.0b protocol, the team could reliably transmit DR signals from a virtual top node all the way to the end use, to increase or decrease the power used by the facility by up to 30 % during a DR event. This was achieved without adverse consequences on the food products stored in the warehouse. The team found that detailed understanding of heat transfer processes in the warehouse, along with an understanding of the refrigeration control system, was necessary to obtain the response desired. The team also found that the desired response with little or no extra energy cost with correct control design. Conversely, an inaccurately tuned control action produces unreliable response and high energy cost. This experience indicates that practical implementation of flexibility measures is best achieved by integrating flexibility with existing cost optimization services, that already use detailed modeling of the facility to minimize operating costs. Beyond storing thermal energy in the product warehouses, other DR opportunities exist through control of resistance floor heaters and charging electric forklifts used to move product within the facility.

Introduction

Globally, the industrial refrigeration sector is dominated by India, the United States and China, with 150M m^3, 131M m^3 and 105M m^3 of the collective 616M m^3 global warehouse capacity respectively (Salin 2018). Warehouses in the United States are overall larger in size but fewer in number, while warehouses in India are smaller in average size but more numerous. This is significant because larger warehouses, on average, use less energy per unit volume per year, as reported in a survey conducted by University of California Davis [Singh 2008]. The annual growth rate of consumer spending on grocery retail varies widely across the world, around 2 %–4 % for a large fraction of developed countries, with a corresponding associated growth rate in logistics infrastructure. In a study prepared by the U.S. Department of Agriculture, there are over 11M m^3 of gross refrigerated warehouse space in the state of California alone (USDA 2017). Assuming an average of 2100 tons of refrigeration capacity [7385 kW_{th}] per M m^3 (Stoecker 1998) and an

average of 1.3–2 kW of electric power per ton of refrigeration capacity (0.37–0.57 kW_e/kW_{th}), this translates to approximately 50 MW_e of total refrigeration electric demand for the refrigerated warehouse inventory in California, where this study focuses. Furthermore, this electric demand represents only a fraction of the total industrial refrigeration load that also includes food and agricultural processors as well as beverage producers (ASWB 2018).

In a study addressing the refrigerated warehouse sector in California (Scott *et al.* 2015), a total potential of 22 MW is estimated. The study also provides detail on the type of warehouses that are most suitable for DR support, notably those that have advanced control capabilities. However, rapidly increasing uptake of online grocery shopping, that translates to substantial increase in demand for frozen and refrigerated goods (reinforced by the Covid-19 pandemic that began in late 2019) combined with technological changes that allow for more sophisticated control of loads, allowing for increased flexibility in managing thermal and other electrical loads at refrigerated warehouses, could see this potential multiplied severalfold (Amarnath *et al.* 2022). A recent study for UN Environment OzonAction (Evans 2018) provides details on individual components of the food chain, providing insight into recent trends in refrigerated warehouse technology, noting that there is significant variation and that efficiency measures could reduce loads substantially. These efficiency measures should be borne in mind when estimating the future of demand flexibility in the refrigerated warehouse sector.

A detailed and comprehensive report written for the Electric Power Research Institute (ASWB 2018) provides an overview of the current state of DR potential of refrigerated warehouses in California, considering emerging needs of the utility sector for flexible loads due to a changing generation mixed that relies increasingly of renewables. The report outlines the variety of grid services that could be provided by the industrial refrigeration sector, and presents the case for standardizing communications between the grid and the customer via OpenADR2.0b, an open-source information model that can transfer signals bi-directionally, for example a DR trigger signal from a load aggregator to a customer, or a reporting signal from the customer to the load aggregator to relate the amount of response that actually took place at the site. From a practical perspective, the presence of a standardized framework to participate in DR lowers barriers to entering the DR market. From a policy perspective, the OpenADR standard was introduced in 2009 and included in the Smart Grid Interoperability Standards Framework, and Federal Energy Regulatory Commission (FERC) identified OpenADR as a key standard for Demand Response, for end uses ranging from residential appliances to industrial loads. More recently, the California Energy Commission has been promoting the further implementation of standardization in DR standards to accelerate progress toward decarbonization goals.

In 2016, the California Energy Commission awarded funding to EPRI (a not-for-profit organization whose mission is to assist the utility industry to provide safe, clean and affordable electricity) to implement a demonstration of demand response capabilities enabled by the combination of a modern and efficient refrigerated warehouse, a sophisticated control system, and the OpenADR communications protocol. The present paper is a summary of this implementation and the outcomes of the research.

Experimental setup

For this research project, the EPRI team partnered with a large logistics services provider, who hosted the demonstration at one of their warehouses, located in one of the logistics hubs in southern California. The plant, depicted in Figure 1, with a gross floor area of approximately 65,000 m^2, is in Southern California Edison's (SCE) electric service territory, but is a direct access customer, meaning that it purchases energy from a competitive provider, rather than from a regulated electric utility. The warehouse is divided into 11 rooms, including freezer rooms (operating at -18 °C), cooler rooms (-18 °C to 0 °C), dry storage area (over 0 °C) and dock areas (4 °C) where food is briefly stored after being unloaded from the refrigerated trucks, before it is placed in the appropriate cold storage room. While



Figure 1. Aerial view of the 65,000 m_2 facility that hosted the demonstration discussed in this paper. Refrigerated semi-trailers loading and unloading at the dock are clearly visible.

the temperature of food items in the cooler rooms and in the dry storage must be kept to within carefully controlled limits, the vast majority of food stored in the frozen rooms can be over-cooled without damage. Therefore, only the frozen rooms are used as thermal storage in this study.

The refrigerant used at this plant is ammonia (NH_3). NH_3 is a common refrigerant for industrial refrigeration due to its high-cycle efficiency. The plant is served by a total of nine compressors, which in groups serve the low, medium, and high temperature zones. The average daily load profile at the plant for October–December 2020 is shown in Figure 2. While the compressors are the largest individual loads, there are clearly other loads that also play a major role in the total plant electricity consumption. The ability to leverage these for DR will be addressed later. The refrigerated warehouse considered in this study was first commissioned in the 1980s, and is representative of many similar plants statewide and nationwide. In recent years, since the acquisition by its current owners, the plant has received several efficiency upgrades including, notably, variable-speed drives on several of the NH_3compressor motors. Operation of the plant refrigeration system is based on controlling temperature. Specifically, proportional-integral-derivative (PID) controllers track temperature by adjusting the fan speed of evaporators located in each room to provide varying amounts of heat removal. In turn, the refrigerant return temperature at the exit of the evaporator coils is monitored to control the refrigerant flow rate and hence the compressor speed and associated electric power draw. Because of the variable speed control, it is possible to fine-tune the power draw and vary power consumption without cycling compressors on and off. The significance of this, for the purposes of this project, is that adjusting the temperature setpoint of individual rooms is the mechanism by which compressor electrical load can be adjusted to provide DR.

Adjusting temperature setpoints to control power consumption is not a new concept at this plant. The same mechanism is used in a process commonly known as "flywheeling," whereby compressor power is adjusted in response to hourly, daily, or weekly forecasts of energy unit cost. The term "flywheeling" refers to the fact that storing thermal energy is analogous to storing kinetic energy in a flywheel. When energy cost is low, compressor power is increased, and vice-versa. The thermal capacity of the frozen food serves as energy storage. While this mechanism appears simple, adjusting power demand to minimize energy cost while ensuring that food temperature is maintained within limits set by the warehouse customers is a complex optimization process. For the warehouse in this study, optimal scheduling of the temperature setpoints is provided by a third-party service provider that specializes in refrigerated warehouse energy system optimization. Specifically, the optimization process adjusts the temperature setpoints to reduce the temperature in the frozen rooms when electricity is less expensive, so that when electricity cost is high the compressors can operate at lower capacity, without allowing the food temperature to exceed a defined upper limit. A recent implementation of the flywheeling process resulted in energy cost savings that produced a payback time of only a few years. In this paper, the same flywheeling concept is extended so that in addition to taking advantage of time-variable energy costs, the same cost-minimization process can also take advantage of

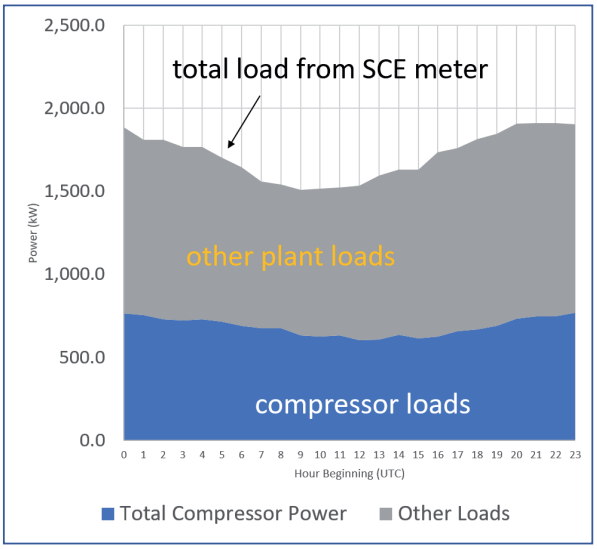


Figure 2. Total and compressor load at the refrigerated warehouse. Note that 19UTC corresponds to noon Pacific Time, meaning that the peak occurs at 2PM Pacific Time.

incentives provided by the utility to adjust demand in times of need. It should also be noted that no such incentives were available in the context of this project, so that adjusting load to support DR was decoupled from the cost-minimization process. In future implementations, DR should be aligned with other cost minimization schemes, for example by adopting transactive energy frameworks, where the full cost of energy, including scarcity and emissions is embedded in bi-directional energy tenders that occur at regular intervals.

The control configuration at the test plant provided an excellent starting point for implementing the DR demonstration. Specifically, to implement DR, it is necessary to adjust the temperature setpoint of one or more rooms – upward to reduce load, and downward to increase load. Moreover, a project requirement was to use the OpenADR 2.0b protocol to adjust the temperature setpoint, and to monitor the response of the plant to the DR signal. The high-level architecture of the refrigeration plant DR reflects the need for the framework to be broadly applicable. For example, it should be able to use legacy system energy management and/or optimization services that may not be able to receive or send OpenADR communications, and different intermediary OpenADR-enabled virtual end nodes that translate OpenADR into a data stream that can be used by the legacy system. The high-level architecture, including the virtual top node (VTN) and intermediary Virtual end Node (VEN) is shown in Figure 3. While this control architecture is still not dominant, scheduling optimization services similar to the one described here are being adopted at an increasing pace, and are likely to be pervasive within the next decade, as existing warehouses are replaced and as mechanical and control systems are upgraded.

The architecture leverages new capability in the OpenADR 2.0b protocol and uses an intermediary to translate OpenADR 2.0b messages to the scheduling optimization system, which is not configured to receive or send OpenADR communications. The use of the intermediary VEN gives the architecture a high degree of flexibility, by allowing legacy systems to interact with the outside world via OpenADR 2.0b. In turn, OpenADR 2.0b

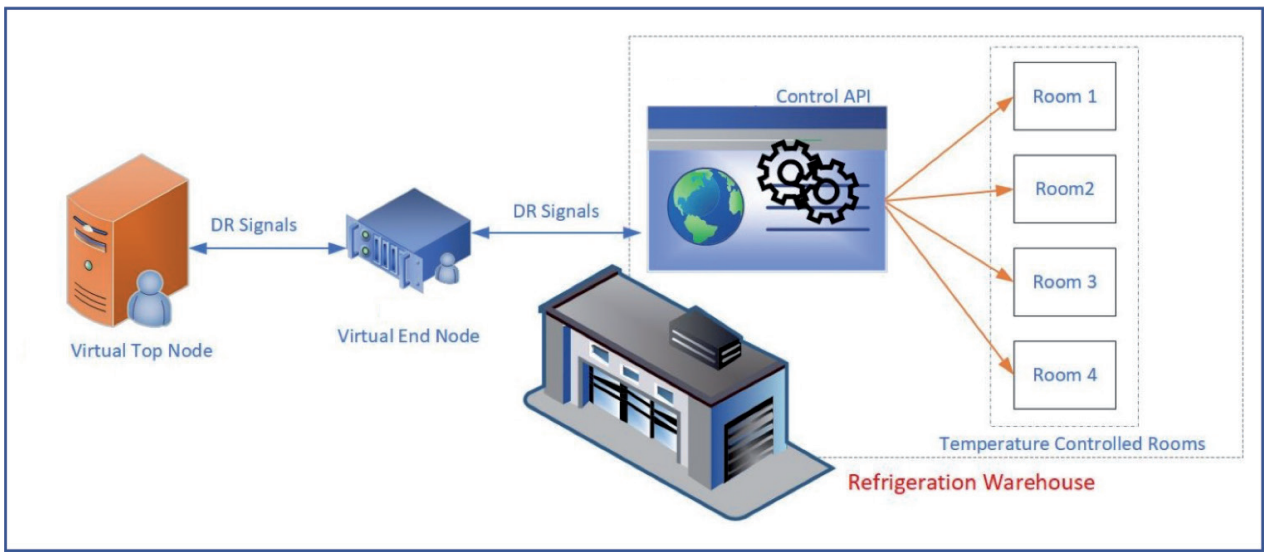


Figure 3. Architecture to deliver a DR signal from a Virtual Top Node to the end-use, via an OpenADR 2.0b intermediary Virtual End Node.

is of interest to the electric utility industry because it allows standardization (and thus, higher adoption), while allowing higher flexibility in relaying information (e.g., by allowing direct reporting of the outcomes of a DR signal). Using this architecture, EPRI (acting as a grid operator or aggregator) uses a VTN to issue a DR notification, which includes information on power level requested and duration via an OpenADR 2.0b message. The message is received by the VEN, and its contents are relayed to the optimization service via a custom Control API. The Control API uses the received information to produce a temperature setpoint schedule and delivers this schedule to the plant controls via an internet-of-things (IoT) server. Submetering information at the plant is accessed by the Control API and sent to the VEN. The VEN interprets the information, and reports this to the VTN at regular intervals via OpenADR 2.0b.

Insights from modeling

To obtain a better understanding of the dynamics of temperature vs. chiller operation in the warehouse frozen chamber, the team constructed a numerical model that captures the essential heat transfer processes, based on these system variables: the heat transfer from the air to the refrigerant at the evaporator, $\dot{Q}_{evap}$; the heat transfer from the food to the air, $\dot{Q}_{food}$; the heat transfer from the warehouse floor to the air, $\dot{Q}_{floor}$; the heat transfer from the warehouse walls to the air, $\dot{Q}_{walls}$. All of these are coupled by a series of simple heat transfer relations, that can be derived readily. The key relations in this model are:

$$\dot{Q}_{evap} = h_{evap} A_{evap} (T_{air} - T_{sat}) \quad 1$$

$$\dot{Q}_{food} = h_{food} A_{food} (T_{food} - T_{air}) \quad 2$$

These equations reveal, quantitatively, that the evaporator transfers heat from the air thereby changing its temperature. Similarly, the temperature of the food, which has a thermal capacity orders of magnitude large than that of the air, can only be changed by the surrounding air at a different temperature. Control of the rate of heat exchange is via the convection coefficient at the evaporator, that can be modeled as:

$$h_{evap}(t) = K_p + K_i \int_0^t \epsilon(\tau) d\tau \quad 3$$

This is simply a proportional-integral-derivative (PID) controller, where K_p is the proportional constant and K_i is the integral constant, with zero derivative term. The forced convection coefficient is bounded so that its maximum value is consistent with forced convection processes in air at heat exchanger surfaces. Natural convection coefficients are estimated based on values suggested in standard heat transfer textbooks. By deriving ancillary. heat transfer relations and obtaining information from the plant operators concerning parameters such as mass of food, total fin area of the evaporator, dimensions of the room, controller constants that produce response similar to the experiments, etc., the team obtained the numerical solution for a heat transfer process that included an initial setpoint temperature drop, followed by a return to default conditions. This is plotted in Figure 4.

As shown in Figure 4(a), from equilibrium initial conditions at -17 °C, the temperature setpoint is reset to -23 °C for the first 10 hours, then returned to -17 °C for the remainder of the time. The temperature of the air drops rapidly for a short time, approximately ½ hour, then follows a slow, essentially linear decline. The temperature of the food, on the other hand, declines linearly for the entire time, following the same rate as the long-term decline of the air temperature. When the setpoint is returned to its original value of -17 °C at hour 10, the air temperature quickly regains several degrees, almost returning to the initial condition, and then begins a linear slow increase. The food temperature inverts its linear decline and begins a slow linear increase.

The explanation for this interesting behavior can be found by considering the evolution of the individual heat transfer processes $\dot{Q}_{evap}$, $\dot{Q}_{food}$, $\dot{Q}_{floor}$ and $\dot{Q}_{walls}$, shown in Figure 4(b). $\dot{Q}_{evap}$ begins at around 375 kW, then rapidly and exponentially declines to a 326 kW steady-state. This is primarily

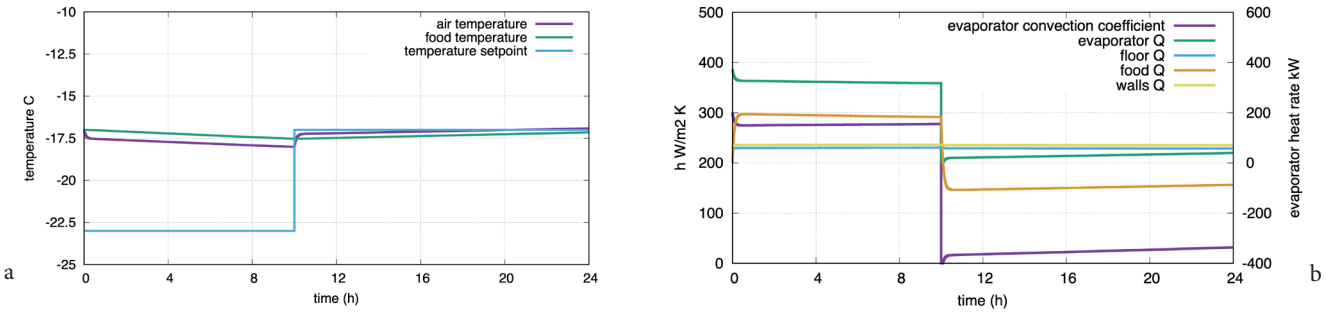


Figure 4. Evolution of frozen room temperatures prior to and following a room setpoint reset (left); evolution of the heat transfer rates between frozen room components prior to and following a room setpoint reset.

a consequence of the control mechanism, whose proportional element is reduced as a consequence of the reduced tracking error, while the integral component increases slowly. At the same time, the food starts to transfer heat to the air as the temperature of the air drops, giving rise to the characteristic slow quasi-linear decline. The floor and wall hear transfer rates do not change appreciably, remaining at around 60–70 kW.

When the setpoint is reset to its original value of -17 °C, at hour 10, the heat flow from the evaporator stops for a short time, while floor and wall heat flows quickly result in an increase of air temperature. After the air becomes warmer than the food temperature, there is a net transfer of heat from the air to the food. As the tracking error reverses sign and becomes much smaller in magnitude, the integral component of the controller becomes dominant and the fans turn on again, at a low capacity, so that the evaporator continues to absorb heat from the air even though the air temperature is slightly lower than the setpoint.

The model revealed important features of the complex heat transfer process dynamics between the evaporators, the freezer room surfaces, the floor heaters, and the food, as follows:

1. The heat transfer process (cooling of the air mass) occurs relatively rapidly, with a time constant on the order of one hour.
2. Once a large enough temperature difference is established (several degrees), effective heat transfer can be established between the air and the food.
3. Because of the high thermal mass of the food (each frozen room contains several thousand tons of food), the rate of change of temperature is slow, even when room evaporators operate at maximum capacity.
4. The system controller tracks the air temperature, not the food temperature, although the heat content is primarily in the food. Consequently, changes in food temperature always lag changes in air temperature.
5. The heat produced by the floor heaters, which eventually ends up in the freezer room space, is likely a significant contributor to the heat balance in the room.
6. Control parameters that are set in the PID control can affect the dynamics of the system, particularly the speed of the compressor, and the activity of the compressor following a temperature reset. Care should be taken in determining the setpoint, because the time integral of the tracking error can lead to unexpected consequences.

Experimental results

EPRI conducted several DR tests in 2020 and 2021, that provided the basis for modelling effort. Based on the modeling insight, the EPRI team triggered further DR events, one of which occurred on the weekend of September 25–26, 2021. This experiment was intended to measure the power response to an initial reduction of temperature setpoint (power consumption should increase), followed by an increase in temperature setpoint (power consumption should decrease). During the weekend, the transfer of food in and out of the frozen rooms is minimal, reducing experimental "noise." As an additional measure to reduce noise, the defrost process for the frozen rooms was de-activated during the DR event. The temperature setpoint and room temperature response, for room 8, one of the frozen rooms, is shown in Figure 5.

The temperature response of room 8 was very close to the analytical model. Prior to the pre-cooling event, the air temperature tracked the setpoint very closely for at least 18 hours. At the onset of pre-cooling, the expected exponential decline of air temperature, lasting less than one hour, was observed, followed by the quasi-linear decline in temperature associated with heat removal from the food. At the end of the pre-cool period, the air temperature climbed exponentially, driven by floor heating, heat transfer from the perimeter walls and roof, introduction of food pallets, forklifts, and defrost cycles. While full validation would require a far more extensive set of experiments than was feasible in the context of this work, as well as much more extensive instrumentation, the outcome of the experiments presented here is consistent with the model and accompanying assumptions. However, additional verification is planned for future projects, that would integrate DR with existing optimization processes.

The corresponding whole-plant power response is shown in Figure 6. Noting that the time is in PDT, the pre-cool period begins at 12 a.m. on September 25, and lasts until 5 p.m. on the same day. The baseline average power draw from all the compressors before the pre-cooling period is 1038 kW. During pre-cooling, the average power draw increases to 1310 kW, 26 % higher than baseline. In the power-down period, follow-

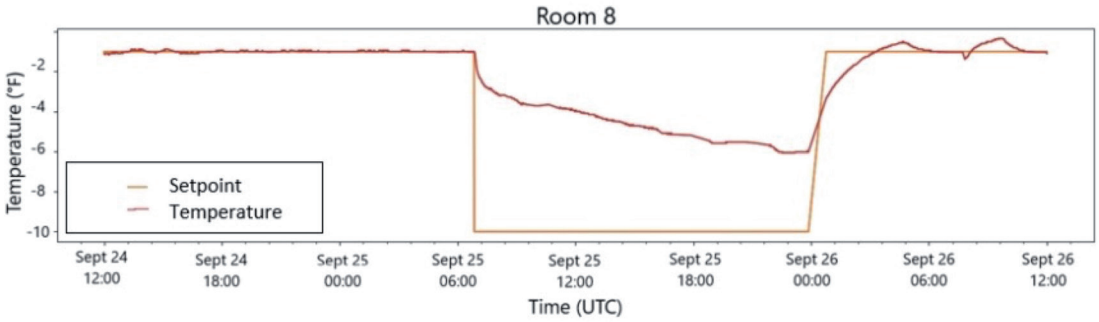


Figure 5. Evolution of air temperature in one of the frozen rooms following setpoint resets.

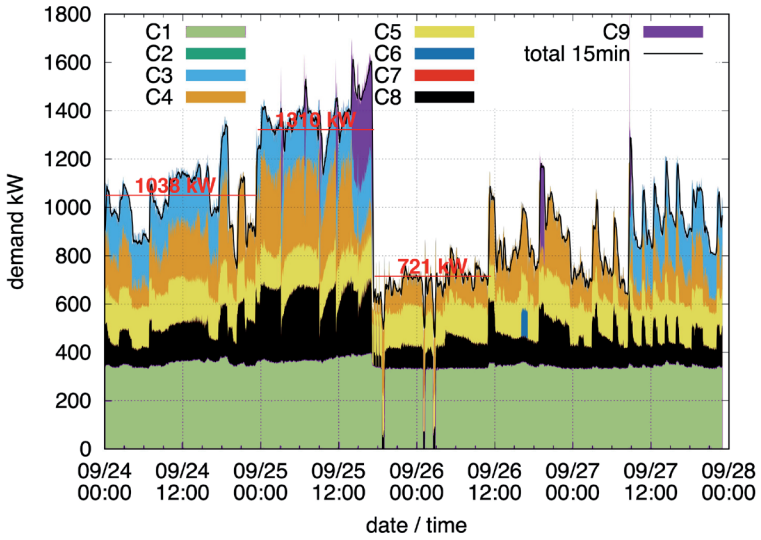


Figure 6. Overall compressor response to the temperature resets in the September 25–26, 2021 experiments. Compressors 3, 8 and 9 are associated with the low-temperature circuit and display the most response.

ing the pre-cooling, power draw drops to 721 kW, a 30 % reduction compared to baseline, and is sustained for a period of the same length as the pre-cooling period, from 5 p.m. on September 25 to noon on the following day. After the DR period, power draw slowly returns to baseline levels.

In addition to the compressor loads, the team identified additional electrical loads that are compatible with DR. These are electric floor heaters, and forklift truck chargers. The floor heaters prevent the ground and foundation of the building from freezing and cracking. While water/glycol floor heaters, using heat recovered from condenser heat rejection, are increasingly used in newer facilities, at this test site all the floor heaters were electric resistance, with combined rated power of 700 kW. The heaters for this plant are either in the "on" state, or in the "off" state, controlled by a simple deadband controller. Nine of the eleven rooms are fitted with floor heaters, but monitoring and automated control is only available in three out of the nine rooms. The team first characterized the operation of the monitored floor heaters by observing their power consumption over the course of one week. The result of this is shown in Figure 7(a). While there is variation, the total heater power on most days is tightly clustered around an average power consumption of approximately 130 kW for the rooms in question, corresponding to approximately 390 kW extrapolated to the entire plant. To test the DR capability of the floor heaters, the team changed the setpoint of the controllable floor heaters from the default value of 1 °C to -18°C, corresponding to the setpoint of the frozen rooms for a two-hour period. With this setpoint reset, it was expected that all heaters would move to the "off" state independently of their prior state. The team conducted a 2-hour-long DR experiment, beginning at 5PM and ending at 7PM. The results of the DR experiment, shown in Figure 7(b). confirm the expectation that all heaters would turn off for the duration of the experiment. Furthermore, this DR action had no measurable effect on the floor temperature. Because the heaters are simple resistances, their response is immediate, and they can be used in conjunction with the compressor DR to improve speed and flexibility, producing *fast and flexible DR*. Forklift chargers are also another potential DR resource; however, this is more difficult to implement because it requires the presence of a forklift charging optimization system to prevent interference with plant operations. The implementation of such a system is currently being considered.

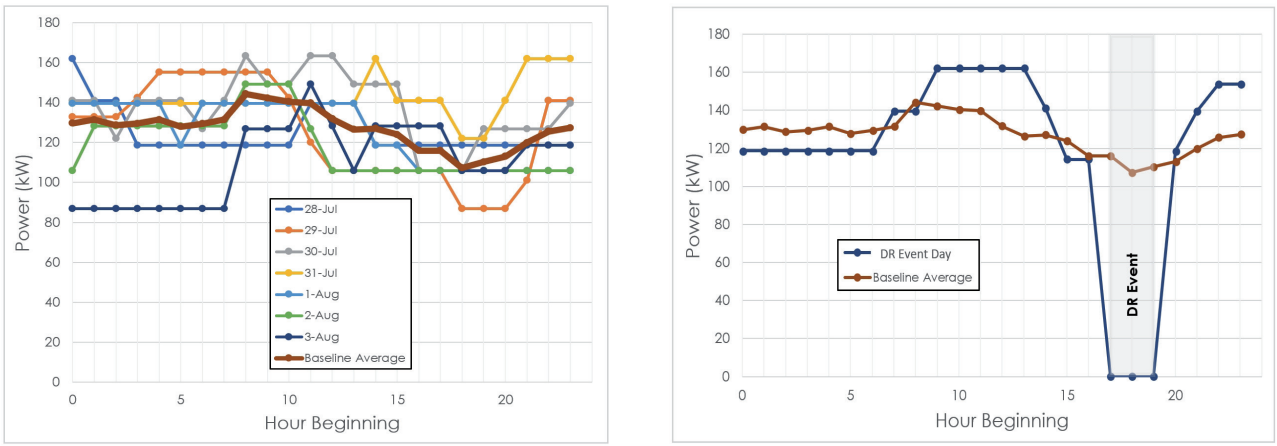


Figure 7. Experiments to characterize DR potential of floor heaters. (a) daily and average combined power draw of floor heaters for three rooms over the course of a week; (b) floor heater power draw on day of DR event, showing that floor heater power is reduced to zero during the 17:00 to 19:00 DR event.

Discussion and conclusions

The main conclusion that can be drawn from this work is that by carefully controlling the temperature setpoints correctly and by setting times that allow the food temperature to decrease, a symmetric power-up (pre-cool) and power-down ratio can be attained from the refrigeration compressors. This enables recovery of the extra energy used to pre-cool the food prior to power-up. This also means that the reduction in power draw during power-down is essentially the same as the increase in power during power-up. This finding is essential because it means that round-trip efficiency is high, and that industrial refrigeration compressors are an excellent DR resource. Moreover, in previous experiments, in which the control of the energy transfer process was less optimized, the team observed substantial round-trip energy losses, meaning that the energy supplied for the food sub-cooling process was considerably more than the energy recovered when the food temperature was subsequently allowed to recover to the normal setpoint. This highlights the need to fully understand the dynamics of the heat transfer and storage processes, and to carefully design the mechanism by which energy is stored in the system. For example, in this work we discovered that transferring energy to thermal storage slowly and at low temperature difference is more efficient than attempting rapid and brief energy transfers.

A secondary and important conclusion can also be drawn. The necessity to pre-cool the food prior to a power-down event complicates planning and thus reduces flexibility. This requirement could be eliminated by making the default setpoint a few degrees lower than the current practice (in this case, 0°F, -18°C). This, in turn, would reduce efficiency for the plant as a whole. On the other hand, the added flexibility may allow increased renewables on the grid and thus overall lower emissions. Similarly, higher costs from reduced efficiency may be offset by higher revenues from participation in DR markets. There is no universal answer to where the default setpoint should be set for optimal flexibility, optimal economics, or optimal emissions. This is a complex problem that can only be solved by integrating flexibility into the legacy optimization system that currently minimizes plant costs, combined with an electricity cost that reflects emissions. Ultimately, the entire cost of providing energy, including the cost of infrastructure, emissions, fuel, O&M should all be included in a single price that is provided in real time and also forecasted at some time horizon. The team also showed that, at the plant level, temperature setpoint can be used effectively to control load indirectly. However, when holistic control of the plant is available, such as the case when plant control is completely integrated with the optimization, it may also become possible to access compressor control directly to further improve the response of the system. Ultimately, this work shows that it is practically realistic to manage large thermal loads to support the grid. Future work, including both scenario simulation and field tests, is needed to fully quantify the potential benefits.

The project demonstrated the effectiveness of end-to-end, bi-directional, automated communication (using OpenADR2.0b). The project team built the end-to-end communication pathway through a combination of "off-the-shelf" components and some "custom" components, mimicking what would be done in a "real-life" project. The team also validated the applicability of the OpenADR 2.0b specification, wherein a return signal is sent back to the DR server validating a demand increase or demand reduction action taken by the end-use system. Such information could be valuable feedback for the distribution system operator, who could then determine how to predict and improve response. Finally, the team identified additional loads that are commonly found at refrigeration facilities, such as floor heating and forklift charging, that could also take part in DR with appropriate system integration. Such additional resources complement the traditional refrigeration compressor loads. Together, these resources have the potential to yield a truly fast and flexible DR resource whose size is considerably larger than current estimates.

References

Amarnath, Ammi, Andrea Mammoli, Angela Chuang, David Showunmi, Don Shirey, Colin Lee, Alekhya Vaddiraj. Steve Hoffman. "Develop and Field Test Flexible Demand Response Control Strategies for Water Pumping Station and Industrial Refrigeration Plant". (2022) CEC Report, Agreement Number: CEC EPC-16-026.

ASWB Engineering and VaCom Technologies. "Analysis of Automated Demand Response Opportunities in Refrigerated Warehouses". (2018). Electric Power Research Institute Technical Report.

Evans, Judith. "Cold Chain Technology Brief: Cold Storage and Refrigerated Warehouses." *OzonAction* (2018).

Salin, Victoria. "GCCA Global Cold Storage Capacity Report." (2018). Texas A&M University Technical Report.

Scott, Doug, Rafael Castillo, Kyle Larson, Brian Dobbs, and Daniel Olsen. "Refrigerated warehouse demand response strategy guide." (2015). Lawrence Berkeley National Laboratory Technical Report.

Singh, R. Paul. "Benchmarking study of the refrigerated warehousing industry sector in California."Public Interest Energy Research (PIER) Program Contract and Research Project Reports. California Energy Commission, PIER Program (2008).

Stoecker, Wilbert F. "Industrial refrigeration handbook". McGraw-Hill Education, 1998.

USDA, National Agricultural Statistics Service. "Capacity of Refrigerated Warehouses 2017 Summary" (2017).