CalNEXT ENERGIZING CALIFORNIA'S FUTURE

Field Study of HVAC Cost Optimized Supply Air Temperature Reset (CORE)

Final Report

ET22SWE0042

Prepared by:

Tharanga Jayarathne,

Rupam Singla, Marc Fountain, TRC

Paul Raftery, Yan Wang, UCB

Curtis Fong, Taylor Engineering

November 11, 2025

Acknowledgements

The authors of this report thank Jared Arwin, John Baffa, and Jim Meacham of Altura Associates for their work on implementation, as well as the building owners and operators, Southern California Edison (funding agency), Energy Solutions (prime contractor), subject matter experts, and other stakeholders for their valuable contributions to this project.

Executive Summary

Multi-zone variable air volume heating, ventilation, and air-conditioning systems are widely used in commercial buildings, but their performance varies significantly. Recent studies have found 4 to 15 percent variation in heating, ventilation, and air-conditioning energy consumption between different supply air temperature reset strategies in multi-zone variable air volume heating, ventilation, and air-conditioning systems. Cost Optimized Reset is a dynamic, cost-responsive, supply-air-temperature-reset solution that minimizes heating, ventilation, and air conditioning cost by optimizing heating, cooling, and fan energy use. The objectives of this research demonstration project were to 1) refine the Cost Optimized Reset algorithm for wider adoption in commercial buildings in California, and 2) demonstrate the operation of Cost Optimized Reset in at least two buildings, evaluating energy and cost impacts.

The team deployed Cost Optimized Reset in four medical office buildings in California within Climate Zones 8, 9, 10, and 14; in three of those climate zones, the team collected and evaluated heating, ventilation, and air-conditioning energy consumption data. We also conducted measurement and verification analysis to compare Cost Optimized Reset with the existing supply air temperature control strategy as a baseline. The results indicated 5 to 22 percent savings in combined heating, cooling, and fan energy; and 1 to 20 percent savings in related energy cost compared to the baseline. The study also uncovered factors in the building and the control system that may have impacted these results, possibly limiting the effectiveness of advanced controls retrofits.

The project team improved upon an existing version of Cost Optimized Reset by adding smoothing factors to improve its operational stability, updating the calculations in the algorithm to account for different types of heating, ventilation, and air-conditioning systems, and creating versions of Cost Optimized Reset that can operate with partial or no zone-level data in multi-zone variable air volume systems so that it may be adopted in buildings that do not have such information available. The team conducted detailed energy simulations to evaluate the improvements, and the different Cost Optimized Reset versions and found them to be consistently better than existing common industry and best practices. The results indicated that Cost Optimized Reset, when compared to existing supply air temperature control strategies, can achieve 1 to 30 percent energy savings in a building with typical loads and airflows in Climate Zones 3 and 7; in some conditions, Cost Optimized Reset can achieve even higher savings. The results also indicated that simplified Cost Optimized Reset algorithms that use less data may achieve comparable performance, increasing the value of the strategy for a wide range of buildings.

Cost Optimized Reset is best paired with a thorough evaluation of the existing heating, ventilation, air-conditioning, and control systems, followed by existing building commissioning, to address any existing issues. Resolving these issues can yield substantial comfort and energy improvements, which are often required to effectively implement zone-demand-based reset methods. The team also identified the importance of supporting the implementers on Cost Optimized Reset deployment in buildings so that its intention is correctly interpreted and programmed. The team believes that Cost Optimized Reset may be better packaged within a programming library offered by building automation system manufacturers so that it is better integrated with the other controls, rather than having implementers program and deploy Cost Optimized Reset as a standalone measure.

Abbreviations and Acronyms

Acronym	Meaning
AHU	Air handling unit
COP	Coefficient of performance
CORE	Cost Optimized Reset
DDC	Direct digital control
FPT	Functional performance test
HVAC	Heating, ventilation, and air conditioning
IOU	Investor-owned utility
IPMVP	International Performance Measurement and Verification Protocol
kWh	Kilowatt-hour
M&V	Measurement and verification
OAT	Outside air temperature
PG&E	Pacific Gas & Electric
RTU	Rooftop unit
SAT	Supply air temperature
T&R	Trim and respond
TOU	Time of use
TMY3	Typical meteorological year three
TOWT	Time of week and temperature
VAV	Variable air volume

Contents

Tables

Figures

Introduction

Multi-zone variable air volume (VAV) heating, ventilation, and air-conditioning (HVAC) systems are commonly used in commercial buildings, but their performance varies widely, leading to a significant performance gap between best practice and typical operation. Two studies led by Lawrence Berkeley National Laboratory (Granderson, et al. 2018) and University of California Berkeley (Raftery, et al. 2018) reported 4 to 15 percent HVAC energy-consumption variation between different supply air temperature (SAT) reset strategies commonly used in VAV systems. Cost Optimized Reset (CORE) is a cost-responsive SAT reset solution for multi-zone VAV systems (Raftery, et al. 2018) that dynamically adjusts the SAT to achieve the lowest HVAC cost considering heating, cooling, and fan energy use, as well as associated electricity and natural gas costs.

The project aimed to deploy and operate CORE in two buildings and evaluate the energy and cost impacts. The project team demonstrated its operation in three buildings with multi-zone VAV systems, and moreover, improved CORE for wider adoption in the state of California. We collected energy consumption data using the existing metering equipment at the demonstration sites and evaluated energy use for CORE against a baseline.

TRC led the project team with the support of University of California Berkeley Center for Built Environment and Taylor Engineers. Altura Associates supported the project team with site recruitment, implementation, and commissioning of CORE, and supplied the project team with data for measurement and verification (M&V).

In this report, the team provides background for the demonstration study, including a discussion of existing SAT control strategies, project objectives, methods and approaches used for energy simulations, field demonstrations, energy consumption evaluation, and findings from the study. The Recommendations section summarizes the team's lessons learned.

Background

Existing SAT reset strategies—including conventional strategies and industry best practice of American Society of Heating, Refrigerating, and Air-Conditioning Engineers (ASHRAE) Guideline 36—determine an SAT setpoint based on factors such as outdoor air temperature and zone heating and cooling requests. These common SAT reset strategies have three inherent deficiencies:

1. They include simplifications and assumptions about the relationship between SAT and total HVAC energy cost.
2. They require tuning key parameters with optimal values that differ for every building and vary over the building's life.
3. There is no easy way to determine what those optimal settings are and whether tuning is improving savings or not.

These limitations motivated the development of CORE, a SAT reset strategy that dynamically optimizes the SAT based on the HVAC energy cost.

The project team developed an initial CORE sequence of operation and implemented it as a test pilot at a University of California, Berkeley's campus office building (Raftery, et al. 2018). This study showed promising results, with 17 percent HVAC energy cost savings for CORE compared to the existing SAT reset strategy of trim and respond (T&R) based on the warmest zone demand. Limited simulation modeling conducted in parallel confirmed the savings to be between 15 and 30 percent over a range of conditions in the same climate. The project team also identified some limitations of CORE, such as compressor or chilled water valve cycling, which occurred when the system was operating near a SAT setpoint that, if slightly lower, would begin to require mechanical cooling. Another limitation was CORE's inability to adapt to older buildings that do not have all the required data, e.g., zone-level data.

Under the current project, the team implemented CORE in three demonstration sites within three different California climate zones and refined the algorithm by introducing features to address key issues identified in the previous study.

Conventional SAT Control Methods

In early implementations of multi-zone VAV HVAC systems, control system designers used fixed setpoints for the SAT, such as 55°F. While they may have chosen the fixed SAT setpoint based on the climate and building thermal load and taken the relative humidity in the outside air into consideration, many designers would likely have defaulted to a fixed 55°F SAT. If outside air has high relative humidity, cooling the air to low SATs ensures that excess moisture is removed at the air handling unit (AHU) and not supplied to conditioned zones, thus avoiding mold and comfort issues related to high relative humidity.

Since 1992, California Title 24 has required SAT reset, with a minimum reset range of 25 percent of the difference between the design SAT and design room air temperature. With this requirement, for a design room air temperature of 75°F, and a design SAT of 55°F, the SAT would have to be able to reset up to at least 60°F, with a common reset up to 65°F. Historically, the easiest way to achieve this was to reset the SAT based on outside air temperature (OAT), where the system increases the SAT when OAT is cooler. California Title 24-2022 prescriptively requires multi-zone HVAC systems to be able to automatically reset SAT based on building loads—i.e., heating and cooling requirements—or OAT. In more recent building automation systems (BAS) where zone-level data are available, zone-demand-based SAT resets have become more frequent.

Through previous research efforts, including interviews with practitioners and reviewing published literature, the project team identified several commonly used SAT reset strategies used in buildings, along with their limitations:

- **Fixed SAT setpoint (no reset):** This method is most often used in older buildings and climates with hot, humid summers. In this strategy, there is no zone feedback, so to avoid not satisfying zone demand, implementations often have an unnecessarily low SAT, which wastes cooling and reheat energy.
- **Outside-air-temperature-based SAT reset:** The SAT is reset based on an OAT range where the SAT is increased when outside air becomes cooler. This method also does not use zone feedback and has the risk of not meeting zone demand or unnecessarily wasting cooling and reheat energy.

- **Warmest zone-based SAT reset:** The SAT is decreased based on the zone with the most critical demand for cooling. T&R logic may be used for the reset. This strategy attempts to meet zone demand but does not balance between the fan, heating, and cooling energy to optimize either the energy use or the energy cost.

In 2018, ASHRAE published the initial version of Guideline 36, which gives standardized sequences of operation for HVAC systems and includes the best-practice SAT reset strategy discussed in the next section.

These conventional SAT reset strategies, as well as ASHRAE Guideline 36, do not explicitly optimize the SAT to reduce HVAC energy or energy cost. Changing the SAT upwards or downwards can either increase or decrease the total HVAC energy and energy cost due to increased or decreased demand for flow rate to meet the zone demand. Therefore, a dynamic reset strategy is required to optimize SAT to reduce HVAC energy and cost.

ASHRAE Guideline 36 SAT Control Method

ASHRAE Guideline 36, titled High-Performance Sequences of Operation for HVAC Systems (ASHRAE 2024), establishes current industry best-practice standardized sequences of operations for HVAC systems. The guideline's focus is to maximize the energy efficiency and overall performance of HVAC systems, provide control stability, and allow for real-time fault detection and diagnostics. The sequences provided in this guideline complement other ASHRAE standards and help maintain occupant comfort and acceptable indoor air quality. By providing uniform sequences of operation that include reset strategies based on real-time building HVAC parameters, ASHRAE Guideline 36 helps in achieving substantial energy savings.

ASHRAE Guideline 36, Section 5.16 provides the standardized sequence of operation of multi-zone VAV air handling units (AHU), including SAT control in 5.16.2. It performs a two-factor SAT reset based on zone cooling demand and OAT and specifies a default OAT range of 70°F to 60°F to reset the SAT upwards from a low limit to maximize economizing. Figure 1 shows an example of a SAT control diagram based on this OAT range.

When the OAT is low, the logic maximizes free cooling by keeping the SAT setpoint as high as possible while still meeting zone demand. As the OAT rises and the economizer cooling is less favorable, the maximum SAT setpoint is reduced. The value of T-Max—represented by the dashed line—is reset between the minimum and maximum cooling SAT setpoints determined by the designer based on real-time demand, i.e., SAT reset requests from zones. This ensures that all zones can get sufficient cold air. During Occupied Mode, this sequence resets the SAT using T&R logic from the minimum cooling SAT setpoint when the OAT is at the maximum limit—and above up to the maximum SAT limit when the OAT is at the minimum limit—by using the minimum and maximum SAT setpoints determined by the designer.

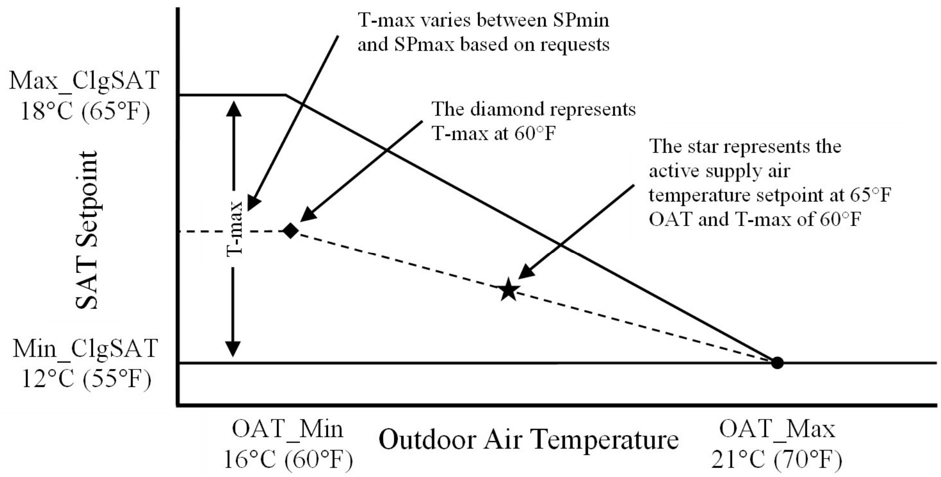

Figure 1: Example SAT reset diagram.

Source: ASHRAE Guideline 36 (ASHRAE 2024)

Implementing ASHRAE Guideline 36 in a building requires the designer to determine the SAT limits, OAT limits, and the parameters required for the T&R reset logic. The SAT and OAT ranges for the reset require tuning for specific building load and occupancy conditions; ASHRAE Guideline 36 provides guidance, but the designer decides.

Initial CORE SAT Control Method

The team developed an initial version of CORE in a research project funded by the California Energy Commission (CEC PIR-12-026) (Raftery, et al. 2018). In multi-zone VAV HVAC systems, resetting SAT impacts heating energy, cooling energy, and fan energy required to meet the zone-level demand. Due to the interactive nature of these different energy uses, SAT reset upwards does not always decrease the total energy use or energy cost. For example, when a zone requires less cooling, resetting SAT upwards can reduce cooling energy use and reheating energy use. However, the airflow rate required to meet the zone demand increases and causes increased fan energy use. The total energy use and the total energy cost depend on the efficiency of the equipment, the tradeoff between the three energy uses, and the energy utility rates. The purpose of CORE is to optimize the HVAC energy cost while meeting the cooling requirements in the building.

CORE does this by first checking if there are cooling requests from the zones. If there are cooling requests beyond a specified threshold, it will reset the SAT downwards until the cooling demand is met, using T&R logic. If no cooling requests are present, CORE dynamically calculates the HVAC energy cost for the current SAT and compares it with an estimated HVAC cost for a small increment and decrement of the SAT from the current setpoint; CORE then resets the setpoint to the lowest HVAC cost out of the three potential SAT setpoints. The CORE algorithm calculates the HVAC energy cost by estimating the heating, cooling, and fan energy uses for the three potential SAT setpoints and the predefined time-of-use (TOU) utility rates for the building, which are input by the designer or user. Figure 2 below illustrates the CORE sequence.

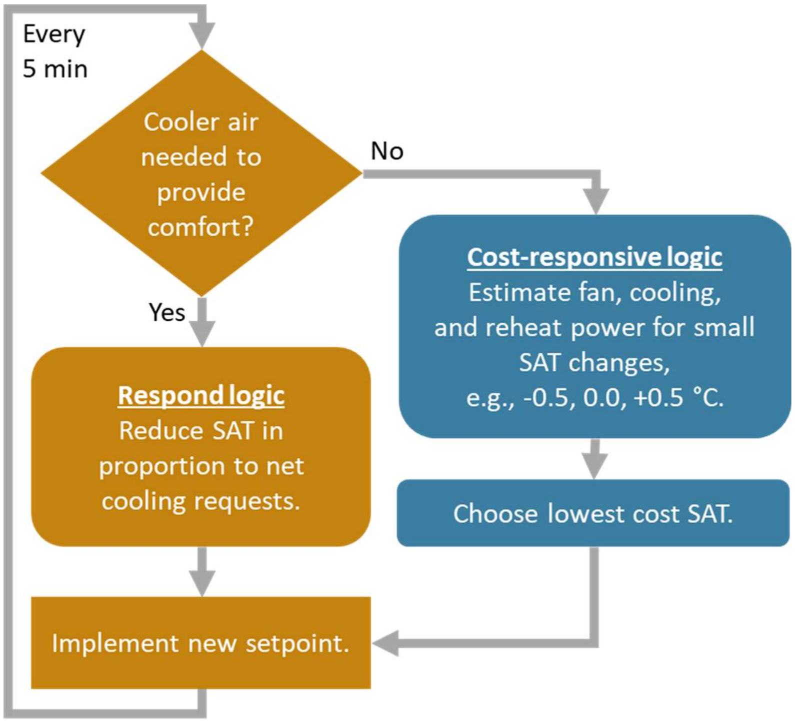

Figure 2: CORE sequence overview.

The initial version of CORE has several advantages over common current practices and ASHRAE Guideline 36:

- CORE dynamically optimizes the SAT to reduce HVAC energy cost compared to common conventional SAT control strategies, such as fixed SAT setpoint and warmest zone-based SAT reset.
- CORE adapts to different building conditions, such as different zone demand levels and utility rate structures, to minimize cost while maintaining comfort. No tuning is required. The algorithm self-corrects over time to set the optimal SAT.
- CORE can improve thermal comfort by reducing overcooling relative to fixed SAT strategy and strategies that do not use zone-level feedback.
- Even though this initial version of CORE reset SAT is based on reducing energy cost, users may choose to reduce other parameters, such as energy consumption or greenhouse gas emissions.

This implementation of CORE in a test pilot site in Berkeley, California, showed 17 percent HVAC energy cost savings compared to the existing SAT reset baseline of warmest zone-based reset with T&R. Under the same study, the team also performed energy simulations to compare CORE with other strategies, including fixed SAT, warmest zone-based reset, and ASHRAE Guideline 36. It used parametric simulations to evaluate the range of savings for various SAT reset parameter values and building load levels. The results suggested that CORE can save 15 to 30 percent more energy than other SAT control strategies.

After initial CORE implementation, the project team identified possible algorithm improvements for testing in this CalNEXT project:

- Using zone-level actual airflow rates led to instability of the algorithm due to airflow control limitations in the terminal units. The team found that this instability could be eliminated by adding a feature that uses the zone-level airflow setpoints instead of the actual airflow for the calculation.
- At certain SAT levels, where a slightly cooler SAT setpoint would require initiating mechanical cooling, the chilled water valve/direct expansion (DX) compressor tends to cycle, leading to unstable SAT control. The team identified that this could be reduced by adding a feature that introduces smoothing for the cooling power calculation.
- Depending on the BAS and its bandwidth, the CORE algorithm operation—which uses zone-level data—can lead to network traffic and system memory issues. This may especially impact existing buildings with older BAS systems where network traffic is a concern. In addition, some sensors, such as zone discharge air temperature sensors, may not be installed in older existing buildings but were required by the initial algorithm, and would therefore restrict the adoption of CORE to buildings that do have these sensors. The team found that a simplified version of CORE could be developed that did not use this zone-level parameter, reducing the number of sensors required along with network traffic load, thereby increasing the scalability of CORE to a wider range of existing buildings.

As described in the Final CORE Algorithm and New Features section, the project team built on the initial version of CORE and incorporated learnings from the pilot implementation to improve operational performance by adding the features and additional versions described above. The team used energy modeling to evaluate the performance of different improved CORE algorithms and compare them with other common SAT reset strategies to inform the current field demonstration.

The project team also developed a version of the CORE algorithm suitable for the state of New York's climate—cold winters and hot, humid summers—under a parallel project funded by NYSERDA (Cost Optimized Reset: Self Tuning Setpoints for VAV Systems). For New York, the team identified the need to account for relative humidity in CORE, along with cooling cost calculation smoothing, zone airflow setpoint use, and the development of simplified versions of CORE to accommodate a wider range of buildings. This New York version of CORE limits the upper end of the SAT based on the outside air dew point to limit the amount of moisture brought into the building. This type of control is not as important in California as it is in New York, so for the version of CORE deployed in this project, the team did not use SAT high limit reset based on the outside air dew point.

Objectives

The objectives of this research were:

1. **Refine the CORE algorithm for broad applicability to existing buildings in California.** The team evaluated additional features that improve the performance of CORE compared to the initial version of CORE in California climates. To improve market scalability, the team investigated different versions of CORE that can be implemented even with BAS limitations, such as when

zone-level data is not available, and evaluated building system requirements, implementer requirements, feasibility criteria, and product-offering requirements.

2. **Implement CORE algorithm in at least two buildings and evaluate energy cost savings.** The team implemented the refined CORE algorithm in four buildings and, in three of them, compared the energy consumption and energy cost under CORE against baseline conditions selected based on industry common practices and best practices.

Methodology and Approach

The CORE Algorithm Improvement subsection describes the methods and approach to meet the first objective, refining the algorithm. The subsections Test Sites Identification, CORE Implementation and Commissioning, and Measurement and Verification Approach describe the methods and approach to meet the second objective, implementation.

CORE Algorithm Improvement

One of the objectives of this project is to refine CORE to be broadly applicable to existing buildings in California. To understand the requirements for the CORE algorithm, the project team evaluated the state of existing buildings, equipment availability, and capabilities of existing BAS. The team also used learnings from the previous research study to add new features that improve operational stability of CORE.

CORE Algorithm Versions

In the previous pilot demonstration (Raftery, et al. 2018), the team observed instability of SAT setpoint control under conditions where the SAT approaches the current outdoor temperature. Here, a small amount of mechanical cooling, like chilled water coils served by the building's chilled water plant, is needed to further reduce the SAT, and the simple algorithm often oscillates around this setpoint. In DX air handling units, this could cause frequent cycling of the AHU's compressors, which reduces equipment life and reduces the stability and performance of the CORE algorithm operation. The team evaluated options to reduce this instability, including introducing smoothing for calculations.

The team also identified that some buildings would not have sensor availability, e.g., zone-level discharge air temperature sensors, or the network bandwidth to implement the initial version of CORE. This led the team to develop simplified versions of CORE that do not use some of or all the zone-level data for SAT reset. The team developed several versions of the CORE algorithm, summarized in Table 1. The Findings section discusses the evaluation of these different strategies through energy simulations, and reports which final version of the algorithm the team decided to use in the field demonstration.

Table 1: Summary of CORE algorithm versions and features.

CORE Algorithm Version	Features
Base	Base algorithm used in the previous research study. This uses zone-level data with discharge air temperature (DAT) sensors. Not used in this project.
CORE 1	To base algorithm, add cooling cost calculation smoothing. Use zone-level data.
CORE 2	To base algorithm, add cooling cost calculation smoothing, remove zone DAT data use. Use zone-level airflow and reheat coil valve position.

Table 2 summarizes the data requirement for each newly developed version of CORE, and ASHRAE Guideline 36 for comparison.

Table 2: Comparison of data requirements for each newly developed CORE algorithm and ASHRAE Guideline 36 SAT controls.

Category	Variable	ASHRAE Guideline 36	CORE 1	CORE 2
Information by designer	Outdoor air temperature bounds	✓		
	Min and max zone airflows		✓	✓
	Min and max supply air temperature	✓	✓	✓
	Trim and response rate	✓	✓	✓
	Utility rates		✓	✓
	Chiller coefficient of performance (COP)		✓	✓
Zone	Discharge air temperature		✓	
	Room air temperature		✓	
	Reheat coil valve position		✓	✓
	Airflow rate		✓	✓

Category	Variable	ASHRAE Guideline 36	CORE 1	CORE 2
AHU	Outdoor air damper position		✓	✓
	Supply and mixed air temperature		✓	✓
	Cooling and heating coil valve position		✓	✓
	Supply or return fan speed		✓	✓
Weather	Outdoor air temperature	✓		
Cooling requests	Room air temperature cooling setpoint	✓	✓	✓

Energy Simulations

The team used energy modeling to evaluate the performance of different CORE algorithms against ASHRAE Guideline 36 controls and conventional SAT control strategies.

The team developed a single-floor model of a commercial office building, meant to represent a single floor of a high-rise commercial office building. The modeled floor area of 14,100 square feet is shown in Figure 3. Zoning is comprised of small office spaces, large office spaces, conference or meeting rooms, and a central mechanical room. The building is served by a VAV AHU with reheat terminals at each of the 24 zones and has a design capacity of 150 people, representing typical occupant density in office buildings. The ceiling and floor of this one-floor model are connected to other floors of a multistory building and are treated as insulated surfaces with no heat transfer. Building envelope characteristics, such as infiltration and surface thermal conductivity, are specified based on California Title 24-2022.

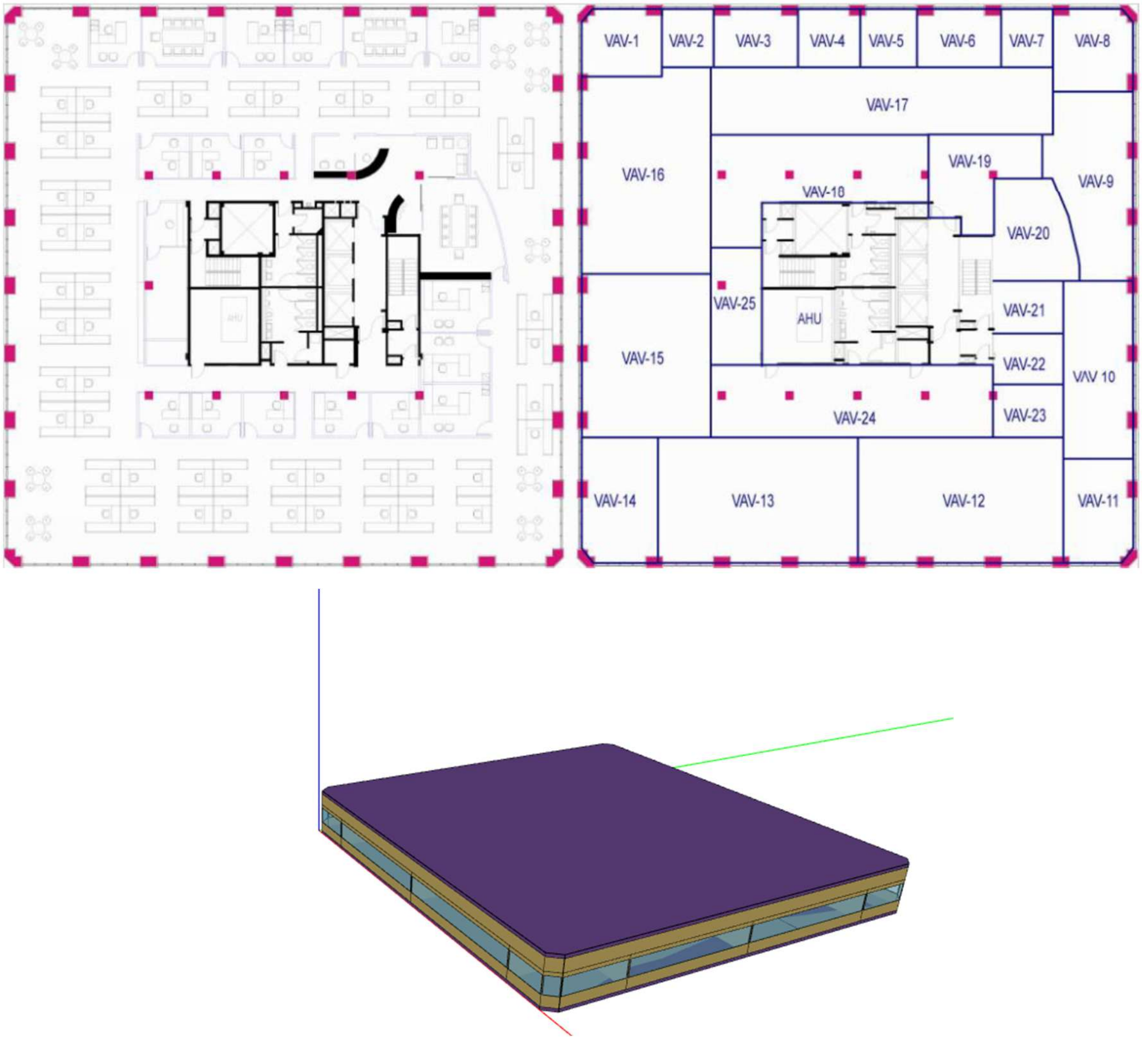

Figure 3: Plan and 3D view of the model building used for simulations of CORE.

The multi-zone VAV HVAC system of the building includes a condensing boiler that uses natural gas to generate hot water for VAV-level reheat and an electric water-cooled chiller to cool air at the AHU level. The central AHU also includes an airside economizer to take advantage of cool outside air to reduce cooling energy.

The team used EnergyPlus, a flagship open-source building energy simulation engine developed by the US Department of Energy. The control algorithms are implemented in EnergyPlus simulations using a recently developed EnergyPlus feature called Python Energy Management System. The Python Energy Management System allows for writing custom control logic in Python and overriding desired simulation variables during the EnergyPlus simulation runtime. The team controlled the parametric variables by modifying the EnergyPlus input data file fields using the Python scripts. They analyzed the results using Python scripts to read simulation outputs in the formats of .csv, .eso, or .sql.

The team used parametric simulations to evaluate how different SAT control strategies perform under different building conditions. Table 3 summarizes the variables the team used for parametric simulations.

Table 3: Parametric simulation variables.

Simulation Parameter	Values
Locations (climate zone)	• Oakland (Climate Zone 3) • San Diego (Climate Zone 7) • Sacramento (Climate Zone 3)
Building type	Single floor of an office building

SAT control baseline strategies	• Fixed SAT:
	  ○ 53°F
	  ○ 55°F
	  ○ 57°F
	• ASHRAE G36 — OA temperature bounds — SAT 55–65°F
	  ○ 50–80°F OAT
	  ○ 60–70°F OAT
	• Warmest-zone-based SAT reset — SAT 55–65°F

Building occupancy schedule	• Commercial office building (baseline) • Educational building (high occupancy)
Internal load schedule	• Baseline • Title-24 ACM (Bucaneg and Wichert 2022) office schedules • 20% of zones empty half the time and 30% of zones empty all the time • Low load: −50% people, lights and plug loads • High load: +30% lighting loads, +125% plug loads

Simulation Parameter	Values
Cooling to fan energy ratio	• Baseline • 25% increase in fan power and 25% increase in rated cooling COP • 25% decrease in fan power and 25% decrease in rated cooling COP
Envelope load	• 25% window-to-wall ratio • 40% window-to-wall ratio (baseline) • 70% window-to-wall ratio
Zone airflow minimum	• Ventilation minimum (baseline) • 20% of design cooling maximum • 30% of design cooling maximum
Utility rates	CA TOU Commercial Pacific Gas & Electric B-10 (PG&E 2024) and G-NR2 (PG&E 2024) rates

Of the above combinations, the team evaluated 12 main parametric simulations for each climate zone with one baseline condition (scenario 1), four scenarios with different internal loads (scenarios 2 through 5), two scenarios with different zone airflow minimums (scenarios 6 and 7), two scenarios of different cooling to fan energy ratio (scenarios 8 and 9), two scenarios of different window to wall ratio (scenarios 10 and 11), and a high-occupancy scenario (scenario 12) presented in theEnergy Model Simulation Results☐ section.

Test Sites Identification

One of the main objectives of this project was to implement the CORE algorithm in at least two buildings in California and evaluate the energy impact against existing, common practice, or industry best practice SAT control strategies. The team developed site-selection criteria and determined that the best way to identify sites that met all criteria would be to find a project partner that could offer demonstration sites through existing customers.

Project Partner and Site Selection

Project Partner: The project team selected the following requirements for a project partner:

- An established supervisory controls contractor operating in a California investor-owned utility (IOU) territory.
- Has a large portfolio of buildings that meet the site-selection criteria discussed below.

- Has an established relationship with building owners and operators.
- Can implement CORE as a supervisory control. This requires the ability to ingest zone airflows, zone discharge air temperatures, zone cooling requests, and write to the SAT setpoint.

Implementing through a supervisory control contractor allowed the partner to program CORE once and apply that same programming across a portfolio of buildings. It also enabled the project team to implement CORE in multiple sites at relatively low cost and allowed the team to validate the supervisory controls contractor as a market adoption pathway.

The project team identified more than 50 potential partners and evaluated them against the criteria above using information available online, prioritizing potential partners with whom the team had an existing connection or had been involved in previous research projects. We then reached out to possible candidates and conducted interviews to understand partner interest, ability to recruit sites suitable for the implementation, and the ability to implement CORE.

The project team selected **Altura Associates**, based in Newport Beach, California, as the partner for site recruitment and CORE implementation.

Site Selection: We developed a detailed list of selection criteria for sites that would meet the CalNEXT program requirement and would allow efficient implementation of CORE. The summarized site-selection criteria are:

- Non-residential building in California IOU electric service area.
- The building is primarily served by one or multiple multi-zone VAV HVAC systems.
- The building HVAC system and controls are in good working condition.
- The project partner can access and provide the project team with all relevant trend data required for CORE implementation and energy use.
- Energy use in the building is sufficiently metered to capture 15-minute interval data for whole-building or HVAC energy use; HVAC energy should be available either as explicitly metered or as a calculated virtual meter. The team also considered using whole-building energy—electricity and natural gas—if HVAC energy could not be explicitly metered or calculated.
- HVAC system trends are tagged following BRICK, Haystack[1], or similar schema. The project partner has validated that HVAC system trends exist and that they are within a reasonable range.
- The building owner is willing to participate in a research project with published data, and the building has no significant planned building retrofits, renovations, or changes in occupancy planned between now and the end of the monitoring period.

Altura evaluated the sites in their existing portfolio against the criteria above and selected seven candidate demonstration sites. The project team and Altura selected three of those buildings in which to start the demonstration, with the remaining buildings being available in case there were issues with one of the original buildings or the team was able to add additional sites. After the implementation in the initial three sites, Altura followed the same procedure to implement CORE in

[1] BRICK (https://brickschema.org/) and Haystack (https://www.project-haystack.org/) are open-source standardized description dictionaries of assets in buildings that provide standardized tagging for all BAS data and control points. This standardization across buildings allows easy adoption of control sequences that are written for such schema.

an additional demonstration site. This demonstration required no guidance from the project team; Altura used their experience and programming from the first three sites. Implementing CORE in additional and more diverse buildings allowed the team to understand performance under different conditions and possible issues and limitations.

Altura provided the project team details of the candidate sites, including building construction and use, HVAC and energy metering infrastructure (whether HVAC energy is submetered), mechanical drawings, and access to BAS trend data via the cloud platform SkySpark[2].

The project team conducted a detailed evaluation to validate that the sites met the recruitment criteria. The team also analyzed the HVAC system and energy meter trends to evaluate data availability and quality.

Demonstration Sites

Table 4 summarizes the CORE demonstration sites, which are all medical office buildings (MOB). The buildings each have a multi-zone VAV reheat system serving most of the building, and all use a Niagara system for HVAC control and SkySpark for trend viewing and storage. The team conducted M&V analysis to evaluate energy and cost savings in MOB-1, MOB-2, and MOB-3. MOB-4 was considered a "light-touch" implementation, where Altura led the entire deployment process with minimum involvement from the research team.

Table 4: Summary of demonstration sites for CORE implementation.

Site	MOB-1	MOB-2	MOB-3	MOB-4*
Building type	Medical office building	Medical office building	Medical office building	Medical office building
Location	Bellflower	Hesperia	Montebello	Wildomar
Building size	33,000 ft^2	49,300 ft^2	51,200 ft^2	83,300 ft^2
Construction year	2023	2021	1988	2024
Hours of operation	M–F: 5 a.m.–11 p.m., Sat: 5 a.m.–5 p.m.	M: 5 a.m.–12 a.m., T-F: 6 a.m.–12 a.m., Sun: 9 a.m.–3 p.m.	M-F: 8 a.m.–6 p.m., Sat, Sun: 7.30 a.m.–4.30 p.m.	M-S: 6 a.m.–8 p.m.
Air handling units	3 rooftop units (RTUs)	3 RTUs	2 AHUs	4 AHUs

[2] SkySpark (https://skyfoundry.com/) is a software platform that is used for building analytics, fault detection, and energy management. It provides real-time cloud access to building BAS trend data for demonstration sites in this project.

Site	MOB-1	MOB-2	MOB-3	MOB-4*
Cooling	DX	DX	Chiller	DX
Heating	Electric reheat	Boiler	Boiler	Electric reheat
Existing SAT control strategy	Warmest-zone-based	Warmest-zone-based	Warmest-zone-based	Warmest-zone- based
Existing SAT range (°F)	55–65	53–58	55–65	55-65
Recent renovations	New construction to an existing site completed in 2023	New construction completed in 2021	Major HVAC renovation completed in 2024	New construction completed in 2024
Large process loads	Very large MRI process loads served by dedicated chillers. These areas were excluded from the analysis	None	None	None

*Light-touch implementation, no M&V analysis conducted.

The team conducted a detailed BAS trend data analysis for the first three sites to identify data availability, data quality issues, and to resolve any issues.

The key results of this evaluation are as follows.

VAV terminal unit trends:

- The team identified four VAV terminal units in MOB-1 that had no or very low airflow rates. Further inquiry confirmed these were shell spaces not currently occupied, which have capped the airflow rate. The team excluded these zones from the implementation at MOB-1.
- At all three sites, the team identified unmapped data points that the team mapped in the BAS. We also identified duplicated data points which were deleted in the BAS.
- The team identified the requirement to detect rogue zones,[3] which would ensure CORE performed well and operated efficiently in the future. To facilitate this, the team added a percentage of hours each VAV terminal unit requested cooling as a metric to the BAS. Building

[3] Rogue zones are zones that consistently demand cooling or high airflow rate and preventing the AHU from effectively adjusting supply air temperature or static air pressure.

operators can monitor this to identify and address if any zone has frequent cooling requests and becomes "rogue."

AHU trends:

- **Mixed air temperature:** This is used in the CORE algorithm but was not measured at MOB-1. Altura subsequently added a calculated mixed air temperature data point using airflow rates and temperatures of supply, return, and outside air. Mixed air temperature was directly measured at the other two sites.
- **Energy metering data**: The team found that HVAC energy use could be calculated from master meter and submeter data. Altura added a virtual meter to do this calculation for the ease of data retrieval at both MOB-1 and MOB-3. HVAC energy was not submetered at MOB-2, and the team used CORE calculated energy for M&V analysis, as described further in Data Collection Plan section.
- **Overall HVAC system trend data:** We identified one zone at MOB-1 that was served by the multi-zone VAV system and a standalone split cooling unit and excluded it from the implementation.

CORE Implementation and Commissioning

The team implemented the refined CORE algorithm in three buildings and estimated savings compared to a baseline. The implementation process involved programming CORE and then commissioning, thorough functional performance testing, and trend data review.

Programming

We shared a plain-language description of CORE with Altura, who then programmed CORE in an offline version of the supervisory building controls module form. This approach allowed Altura to test the CORE operation and fix any issues before implementing the algorithm in the building, after programming and initial testing. With Altura's support, the research team conducted a preliminary review of the programming to ensure CORE logic was correct. Altura then implemented the CORE algorithm in the selected sites' BAS.

Appendix A provides the adapted CORE algorithm implemented at the three medical office buildings, each with a multi-zone VAV system. Unlike many multi-zone VAV systems in California, the MOB-1 site has electric resistance reheat at the VAV terminal units; accordingly, we updated the CORE algorithm to account for electric reheat under the heating cost calculation.

Figure 4 and Figure 5 below show two screenshots of BAS programming at the MOB-1 CORE implementation. Figure 4 shows the zone-level airflow calculation block where the algorithm estimates the airflow rate for the alternate SAT conditions, i.e., ±0.5°F from the current SAT. The airflow rate estimation node is highlighted in red. Figure 5 is a screenshot of the equation used in this node. The airflow rate estimation from all VAV terminal units, highlighted in green in Figure 4, is used to estimate the HVAC energy cost at the system level.

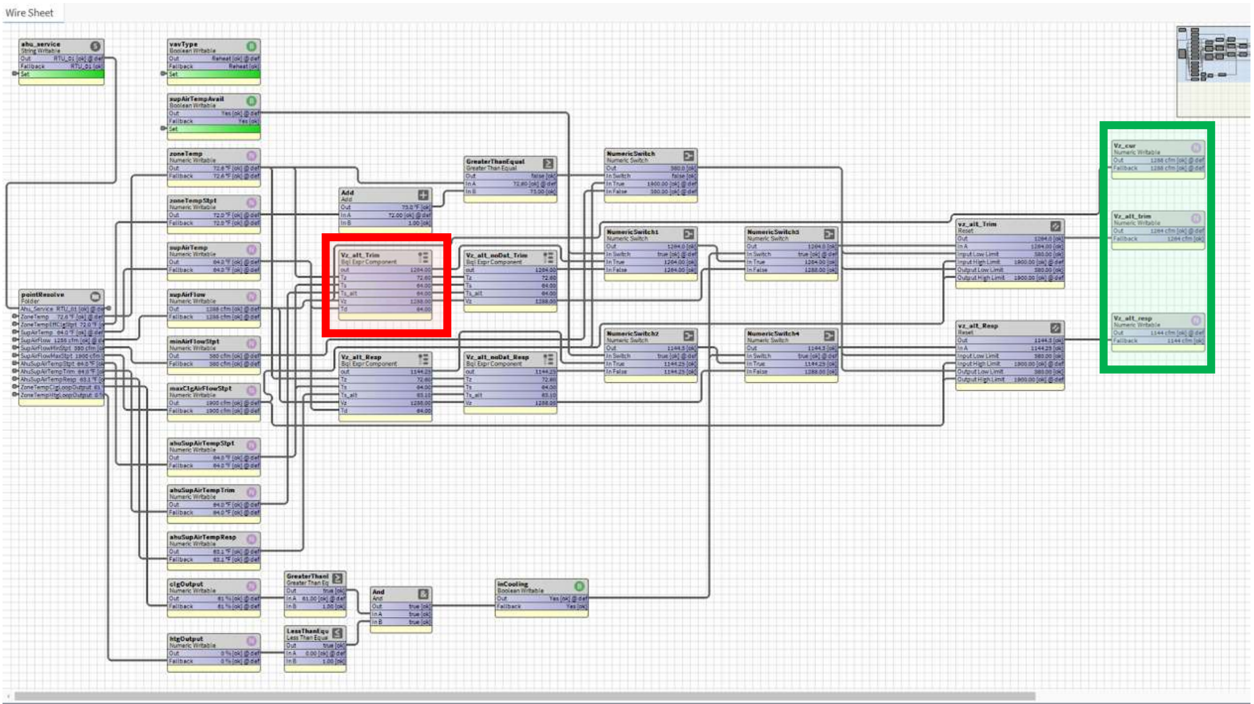

Figure 4: Screenshot of BAS programming showing the zone-level alternate airflow estimates for the T&R.

Note: The green box highlights the T&R outputs; the red box highlights the estimated airflow calculation for an alternate SAT case.

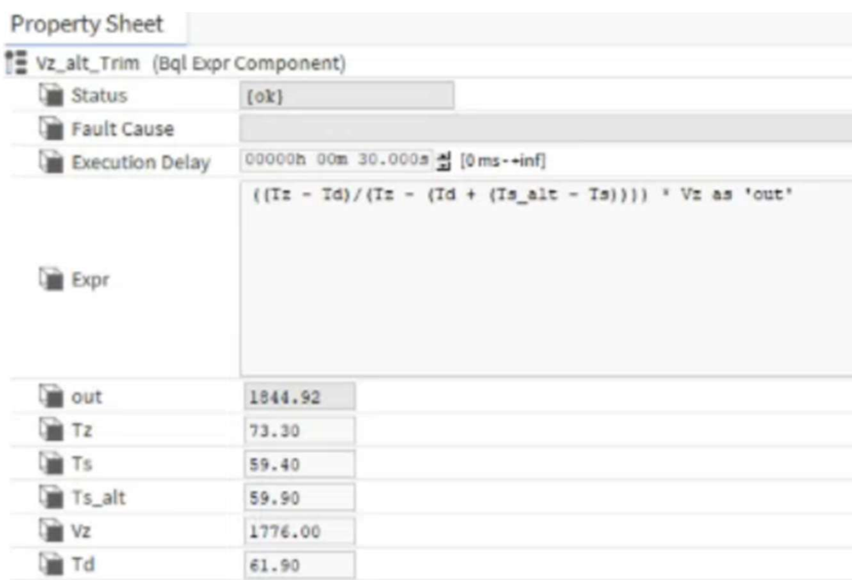

Figure 5: Screenshot of BAS programming showing the airflow calculation equation for an alternate SAT case.

Note: This alternate SAT case is the red box in Figure 4.

Functional Performance Testing

The project team developed functional performance test (FPT) scripts to evaluate SAT reset by the CORE algorithm and shared them with Altura for their use. The test scripts were Microsoft Excel-based with macros, allowing Altura to perform the test as specified, record the relevant data from the BAS to verify performance, and save any screenshots of BAS relevant to the test. The test scripts were accompanied by detailed step-by-step instructions on how to use them and referred to training videos developed by Taylor Engineers.

Altura conducted FPT to systematically test the BAS to ensure that the CORE algorithm was programmed correctly. When conducting FPT, Altura manually overrode building and outside measurements and control points, such as the number of cooling requests, to observe how the BAS behaved under CORE. Through this process, Altura simulated the basic conditions that impacted the CORE calculation, observed the responses, and compared them to expected responses.

The FPTs tested four areas of the BAS:

- **Zone calculations:** FPTs evaluated if zone cooling requests and airflow rate setpoints were calculated correctly by overriding the SAT setpoint at the AHU level and checking if the zone airflow rate setpoint and the reheat coil adjusted as expected.
- **System calculations:** FPTs evaluated if the CORE algorithm calculations of HVAC cost for SAT current setpoint and next step setpoint options (±0.5°F) were correct by overriding the OAT reading and checking if the heating, cooling, and fan cost calculated in the BAS was equal to the expected cost.
- **T&R logic based on the CORE algorithm:** FPTs evaluated if the T&R reset was working by overriding the fan energy cost to zero and checking if the SAT setpoint increased as expected.
- **T&R logic based on cooling requests:** FPTs evaluated if the T&R reset was working by increasing the number of cooling requests and checking if the SAT setpoint decreased as expected as shown on the left side of Figure 2.

The project team reviewed the FPT results and provided feedback on issues that needed to be addressed and identified FPT re-runs that were needed before the team accepted the FPT results.

Trend Data Review

After successful completion of FPTs and updating of the programming, the research team conducted a trend data review using the data available via SkySpark. The team allowed the BAS to run CORE for several weeks before retrieving the data from this period to perform the trend review. The team reviewed how the SAT behaved with other parameters—zone demand, fan energy cost, heating energy cost, and cooling energy cost—to determine whether CORE behaved as expected. To make this assessment, we specifically used trend data of the AHU and rooftop unit (RTU) SAT, number of cooling requests from VAV boxes, supply and return fan energy use, cooling system energy use, and heating system energy use. The team identified several issues and worked with Altura and the building managers to address them or find workarounds. These issues are discussed under the CORE Implementation Efforts and Issues section.

Due to the nature of the issues discovered, the research team also conducted a more detailed review, including intermediate calculations of CORE using AHU and RTU and VAV box temperature, flow rate, requests, and coil operation data. This detailed review uncovered additional issues related to programming and the building systems, which the team resolved. Lessons learned from these issues are also discussed under the CORE Implementation Efforts and Issues section.

Measurement and Verification Approach

Measurement and Verification Methodology

The team conducted measurement and verification (M&V) to meet our objective of evaluating the energy and energy cost savings of CORE compared to a baseline SAT control strategy at each demonstration site. Building conditions and climate conditions affect HVAC energy use and can result in different energy consumption results for different periods, even if the same strategy is used. Therefore, an analysis using a normalizing model was essential to accurately compare different strategies run during different time periods. To estimate energy and energy cost savings, the project team followed established International Performance Measurement and Verification Protocol (IPMVP) procedures to develop normalized regression models for energy use for CORE and baseline strategy energy data.

The team used the IPMVP Option B, "Retrofit Isolation with All Parameter Measurement," to quantify HVAC gas and electricity savings from the CORE implementation, using the following equation, General Equation 1 in Chapter 4.1 of IPMVP protocol:

$$\textit{Savings} = (\textit{Baseline-Period Use} \text{ or Demand} - \textit{Reporting-Period Use} \text{ or Demand}) \pm \text{Adjustments}$$

We used IPMVP Option B Retrofit Isolation as the M&V method in MOB-1 and MOB-3 buildings. In MOB-1 and MOB-3, HVAC electric and natural gas energy consumption were submetered; in MOB-2, the team used the IPMVP Option C Whole Facility method. The CORE algorithm calculates heating, cooling, and fan energy use at five-minute intervals based on airside measurements. Where submetered HVAC energy use was not available, and to supplement it where it was available, the team used the calculated heating, cooling, and fan energies. The former method captures the actual metered energy use, including additional equipment energy use from pumps and cooling towers, and indicates the performance of CORE while taking equipment efficiencies into account. The latter—calculated heating, cooling, and fan energy—specifically captured the energy used for CORE optimization, as Option B isolates the specific energy use of the entire HVAC system for CORE versus baseline SAT control strategy. Isolating HVAC energy use ensures that non-HVAC loads in the building do not affect savings estimates. This is especially important in sites like the MOB-1 medical office building, which has large MRI scanner loads.

The team developed regression models using a time-of-week-temperature (TWOT) model (Mathieu, et al. 2011) for both CORE and baseline energy usage. They compared energy savings based on standardized weather data such as typical meteorological year three (TMY3). At each demonstration site, the team characterized energy savings in terms of both the normalized annual energy use (kWh) and the actual energy use difference and the actual energy use difference—both in kilowatt-hours—during the monitoring period. The Energy and Cost Savings Results section presents results from the data collected to date.

Randomizer Measurement and Verification Sampling Procedure

To capture the full range of weather conditions, M&V on research field studies typically involves running the baseline and intervention conditions in each building for up to nine months each. In randomized M&V, a sequential method of running the baseline and intervention conditions is used to significantly reduce the total M&V period (Raferty, et al. 2024). In this project, the team programmed BAS to randomly select between CORE and baseline strategies at midnight every day and run the selected strategy for the day. Over time, this ensured that the distribution of days where either strategy run had the same range of building operating conditions.

Figure 6 illustrates the randomization schedule the team created for 20 weeks, blocked by day of week. This randomization ensured that both strategies ran roughly equal amounts and covered all types of days of the week to ensure that all events in the building—e.g., weekends, weekly seminars, etc.—were represented in both strategies.

Figure 6: Randomized rapid M&V schedule used at demonstration sites.

Note: Grey cells represent holidays, which were run as baseline periods and excluded from the analysis.

Baseline Conditions

The team evaluated the existing control strategy of each site and decided to use it as the baseline condition for rapid M&V. The existing strategy was similar to ASHRAE Guideline 36, where T&R logic is used to reset the SAT setpoint based on cooling requests from zones. However, the existing strategy did not use OAT to determine the SAT reset range, as prescribed in ASHRAE Guideline 36 and illustrated in Figure 1; instead, it used fixed reset ranges, which are available in Table 4.

Data Collection Plan

Altura provided the project team access to the SkySpark database platform that recorded the BAS trend data for all three buildings, where we could view and download trend data of energy-related parameters—such as HVAC electric energy, whole-building electric energy, and whole-building natural gas usage—as well as temperature and flow rate data. The team reviewed energy data from SkySpark weekly and evaluated it for common system and data quality issues, such as missing data, unexpected values, and equipment failure. We notified and worked with Altura and building management to address any issues identified in these reviews.

Table 5 summarizes the data the team used for energy calculations. The virtual meters for HVAC power in MOB-1 and MOB-3 provided calculated electric power usage from the HVAC system, which used a calculation that subtracts power from submeters for different equipment and lighting systems from the main electric power meter to isolate HVAC power usage. In addition to the data listed here, the team used CORE enable status, occupancy status, and OAT for each RTU and AHU.

Table 5: Summary of data points used for energy calculations.

Building	Parameter	Units	Frequency	Source
MOB-1	HVAC electric power (virtual meter), all 3 RTUs and zones	kW	5 minutes	Calculated from master electricity meter and virtual meters
	CORE heating energy, each zone	kWh	5 minutes	Calculated by CORE algorithm based on airside measurements
	CORE cooling energy, each RTU	kWh	5 minutes	Calculated by CORE algorithm based on airside measurements
	CORE fan energy, each RTU	kWh	5 minutes	Calculated by CORE algorithm based on fan speed percentage
MOB-2	Whole building electricity energy	kWh	5 minutes	Master electricity meter
	Whole building gas energy use	Therms	5 minutes	Master gas meter
	CORE heating energy, each RTU	Btu	5 minutes	Calculated by CORE algorithm based on airside measurements

Building	Parameter	Units	Frequency	Source
	CORE cooling energy, each RTU	kWh	5 minutes	Calculated by CORE algorithm based on airside measurements
	CORE fan energy, each RTU	kWh	5 minutes	Calculated by CORE algorithm based on fan speed percentage
MOB-3	HVAC electric power (virtual meter), both AHUs	kW	15 minutes	Calculated from master electricity meter and virtual meters
	Heating gas energy use	kBtu/h	15 minutes	Calculated from gas meter flow rate
	CORE heating energy, each AHU	Btu	5 minutes	Calculated by CORE algorithm based on airside measurements
	CORE cooling energy, each AHU	kWh	5 minutes	Calculated by CORE algorithm based on airside measurements
	CORE fan energy, each AHU	kWh	5 minutes	Calculated by CORE algorithm based on fan speed percentage

Data Processing

The team reviewed the baseline and CORE energy data collected and removed periods from the analysis with significant data quality issues, and we then evaluated the data for gaps and known issues during the data collection period before deciding whether to remove any periods from the analysis. Specifically, the team removed data from periods where issues in the system or programming were known to have affected the operation or where there were large data gaps of more than one day due to communication issues. For the remaining data, no gaps in energy consumption data were larger than two hours; these gaps were filled in with the last measured data value. We found a few instances where values were unusually high, likely due to erroneous readings, and removed these from the analysis. The team also removed scheduled unoccupied periods from this analysis using the occupied status reported in SkySpark, we and we used cleaned data for energy regression model development.

Energy Savings Analysis

ACTUAL ENERGY USE DIFFERENCE

The team used a one-tailed hypothesis test to compare the measurement distribution of the two sampled control strategies. The null hypothesis is stated as follows: The mean daily energy usage sampled from the intervention measurement set is greater than or equal to the mean daily energy usage sampled from the baseline measurements (H_0: $\mu_{interv} \geq \mu_{base}$). The alternative hypothesis, which describes the target effect, is stated as: The mean daily power usage sampled from the intervention measurement set is less than the mean daily power usage sampled from the baseline measurements (H_a: $\mu_{interv} < \mu_{base}$).

By comparing the test statistics, specifically the p-value, an analyst can evaluate whether there is sufficient evidence to reject the null hypothesis in favor of the alternative. In this context, a low p-value (typically less than 0.05) indicates that the observed energy savings are statistically significant and unlikely to have occurred by random chance. A p-value below the chosen significance level supports rejecting H_0, thereby validating that the intervention from CORE results in lower average energy use compared to the baseline. Conversely, a high p-value suggests that the observed differences may not be statistically distinguishable, and the null hypothesis cannot be confidently rejected.

ESTIMATED NORMALIZED ANNUAL ENERGY CONSUMPTION

The team used the baseline and CORE data to develop separate models using the time-of-week and temperature (TOWT) approach and determined the normalized energy consumption to characterize the annual normalized energy consumption. The TOWT approach uses the time of the week and OAT as the independent variables to develop a changepoint regression model.

The team applied these models to an annual weather file using TMY3 weather data to determine typical annual load profiles, which allowed comparison of the two strategies for the same weather conditions. An example regression equation is shown below.

Equation 1. Sample Regression Equation

$$E_i = \alpha_i + \sum_{j=1}^{6} \beta_j T_{c,j}(t_i) + \sum_{k=1}^{n} \gamma_k S_k$$

Where:

E_i	= Energy consumption at time interval i
t_i	= i^{th} time interval
$T_{c,j}$ (t_i)	= Component temperature computed based on algorithm
α_i	= Coefficient for time of day and week estimated by regression
β_j	= Coefficients for component temperatures estimated by regression
γ_k	= Coefficients for independent variables determined by survey responses

S_k = Independent variables, e.g., number of occupants and thermostat setpoints, determined by survey questions and responses

The total difference between the baseline and CORE energy use is the expected normalized energy savings of CORE and can be expressed using the equations below.

$$Q_{elec,savings} = Q_{elec,baseline} - Q_{elec,CORE} \; [kWh]$$
$$Q_{ng\_htg,savings} = Q_{ng\_htg,baseline} - Q_{ng\_htg,CORE} \; [Btu]$$

The team followed modeling best practices and selected a data collection period to cover a range of weather conditions. We assessed model accuracy using the following targets for sub-hourly data, based on ASHRAE Guideline 14[4] and industry guidance (Granderson, et al. 2019):

- Coefficient of determination (R^2)> 0.75
- Coefficient of variation of the root mean squared error (CV(RMSE)): < 30%
- Normalized mean bias error (NMBE): < +/-10%

The team collected energy and energy cost data, presented in the Findings section, to determine total annualized electricity and gas usage, a monthly breakdown of energy consumption, and average daily energy use profiles for different weather conditions throughout seasons. These results can also indicate if the system behaved as expected and if there are operational issues that need to be addressed or considered under data cleaning. The results can also be compared against similar buildings based on previous studies and publicly available data.

UTILITY COST IMPACT

The sites generally have TOU rate plans for electricity, where the utility charges a fixed price per kWh based on the time of day and time of year, with the late afternoon and evening periods having a higher rate than other times of day and summer season rates higher than those in winter. While electricity peak demand charges can be a large percentage of a total utility bill, considering these charges was outside of the project scope so the team did not consider them.

Utilities generally use a tiered rate—with a fixed price per therm in each tier—for natural gas usage, in which total gas consumption puts customers into one of two or three tiers. For this analysis, the team assigned customers to the appropriate tier and then used the total savings from the retrofit determined in the Findings section to determine utility cost savings. Table 6 below shows utility rate information for each site. In addition to the hourly rates, the team recognized that electricity peak demand charge can be significant for commercial customers, but we did not evaluate the impact on peak demand or the related cost.

[4] ASHRAE Guideline 14-2014 Measurement of Energy, Demand, and Water Savings https://webstore.ansi.org/standards/ashrae/ashraeguideline142014

Table 6: Utility rates at each site.

	MOB-1	MOB-2	MOB-3
Electricity			
Midnight–8 a.m. rate ($/kWh)	0.04199	0.04199	0.042
8 a.m.–4 p.m. rate ($/kWh)	0.03919	0.04199	0.042
4–9 p.m. rate ($/kWh)	0.04229	0.04375	0.042
9 p.m.–midnight rate ($/kWh)	0.04199	0.04199	0.042
Natural gas			
First tier	N/A	0–100 therms	0–250 therms
Second tier	N/A	> 100 therms	> 250 therms
First tier rate ($/therm)	N/A	1.16757	1.125
Second tier rate ($/therm)	N/A	0.96065	0.65998

Source: Altura Associates

Findings

Overview

This section presents the energy model simulation results, the final CORE algorithm and added features, CORE implementation learnings, energy consumption and savings results, and an assessment of CORE's market scalability. In this Findings section, we use the following shortened category labels: CORE 1 and 2, ASHRAE Guideline 36, Baseline, Fixed SAT, Warmest SAT.

Energy Model Simulation Results

Results Summary

Background: The team conducted comprehensive energy model simulation studies to evaluate the performance of the CORE algorithm against other existing and industry best practice SAT control strategies. This section presents results for the baseline internal load schedule, fan and energy ratio, and envelope load specified in Table 3 in California: Oakland, San Diego, and Sacramento. The team conducted parametric simulations to understand the impact of weather on CORE, especially in mild climates that allow more free cooling.

Key Findings: In terms of HVAC energy cost minimization, **CORE 1 and CORE 2 demonstrated the best performance** across all scenarios.

Details: The team ranked the performance of eight control strategies based on the annual HVAC energy cost across 12 Oakland scenarios, as shown in Figure 7, with darker blue shading indicating lower ranked costs and darker red indicating higher ranked costs. The energy cost difference is presented in both annual dollars per square meter per year and the percentage difference relative to the best control strategy in each scenario. As HVAC system design and operations vary significantly across the commercial building sector, it is important for an optimal control strategy to achieve consistent energy savings across scenarios.

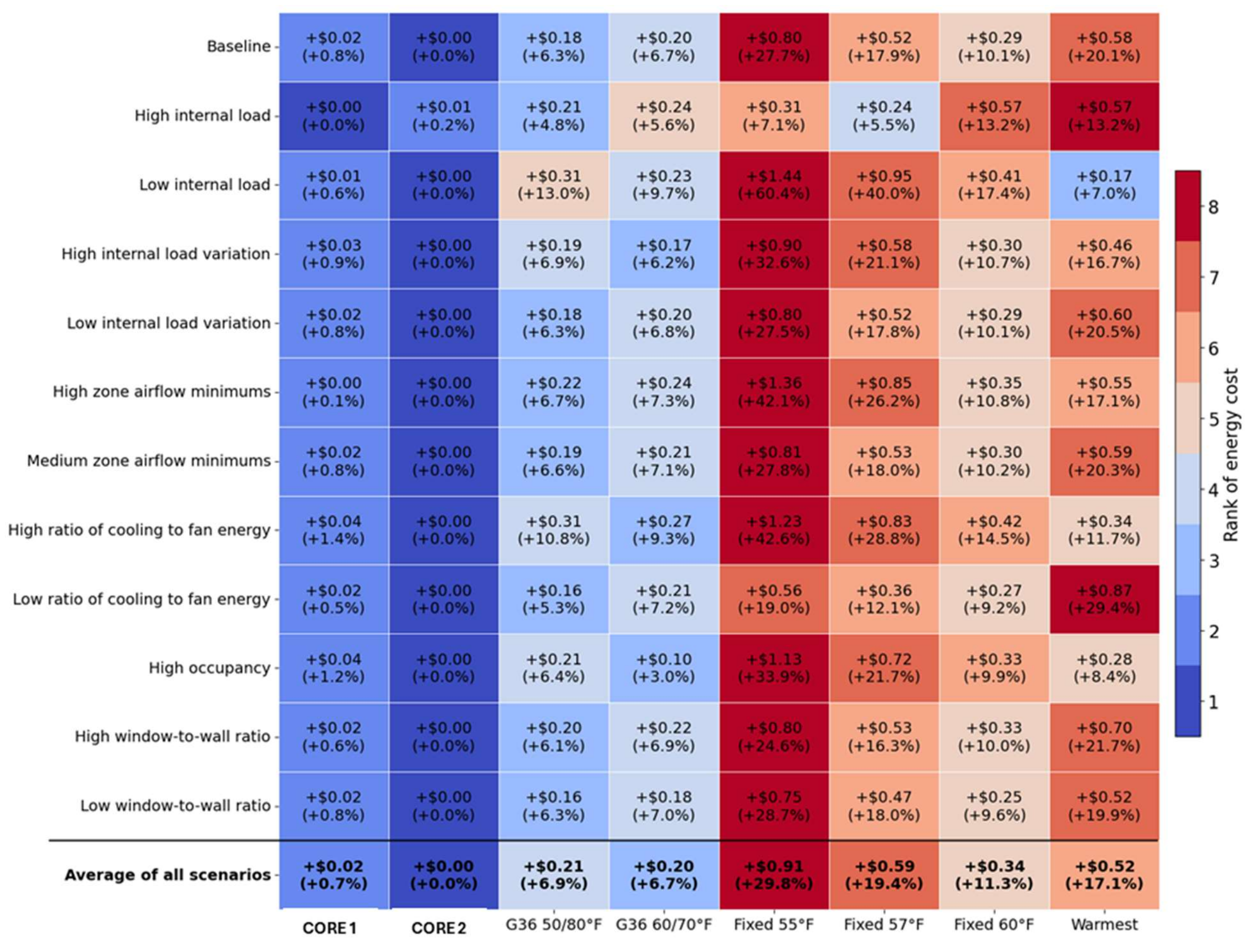

Figure 7: Rank of HVAC energy cost ($/m²· year) for Oakland parametric simulations with TOU utility rate.

Overall, there was a difference in energy cost between more sophisticated control strategies such as CORE and ASHRAE Guideline 36, and simpler ones, like Fixed SAT. On average, compared to Fixed and Warmest SAT control strategies:

- CORE achieved 10 to 30 percent ($0.3–0.9/m²· year) energy cost savings

- ASHRAE G36 achieved 4 to 23 percent ($0.13–0.7/m²· year) energy cost savings

The team also carefully assessed the unmet hours of heating and cooling setpoints, which were consistently low between different control approaches in each scenario. This finding suggests that occupants were comfortable with the temperature of their space, despite the variations in energy costs. Overall, the simulations indicate that CORE consistently achieved the highest energy cost savings in all parametric simulations without compromising the occupant's thermal comfort.

Energy Use Breakdown

To evaluate performance variation between CORE and other control strategies, the team looked at the energy and energy cost breakdown for heating, cooling, fan, and other HVAC energy use, as illustrated in Figure 8. Gas usage for heating was comparable across all strategies, while cooling and fan energy were used differently throughout strategies to meet the demand. Compared to existing strategies, CORE and ASHRAE Guideline 36 balanced fan and cooling energy to achieve lower overall energy and energy cost.

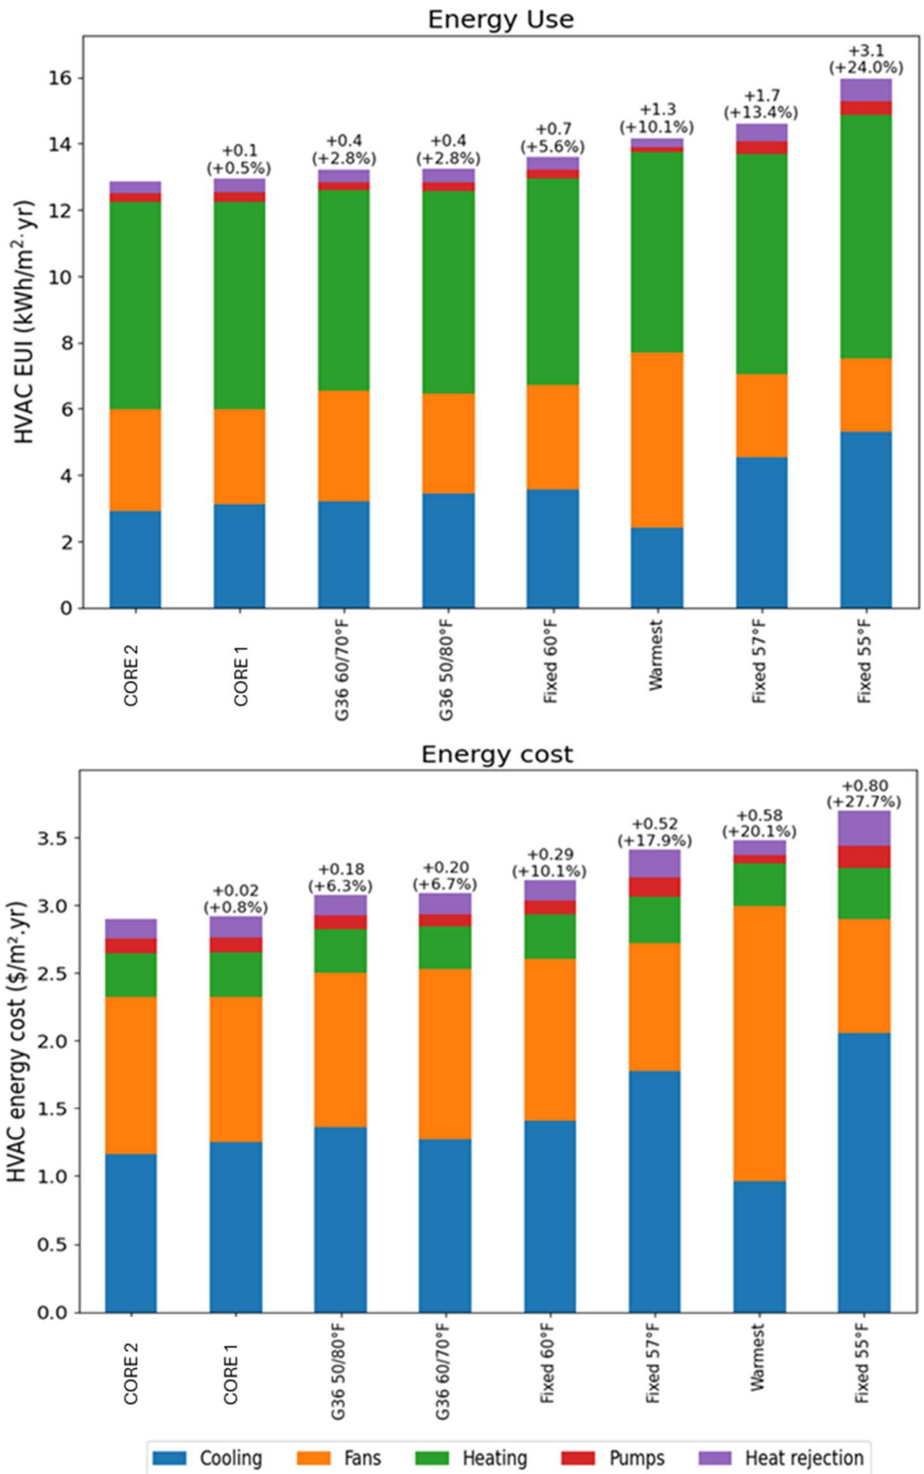

Figure 8: Comparisons of HVAC energy use and cost for the Oakland Baseline (TOU utility rate).

Note: These comparisons are arranged in ascending order according to their respective HVAC energy consumption and costs.

Figure 9 shows the total HVAC energy cost and component breakdown for San Diego. As with Oakland, CORE 1 and CORE 2 had the lowest HVAC energy costs, followed by ASHRAE Guideline 36 and other conventional strategies. Warmest-zone-based SAT reset had the highest cost and was significantly more expensive than the Fixed SAT strategies.

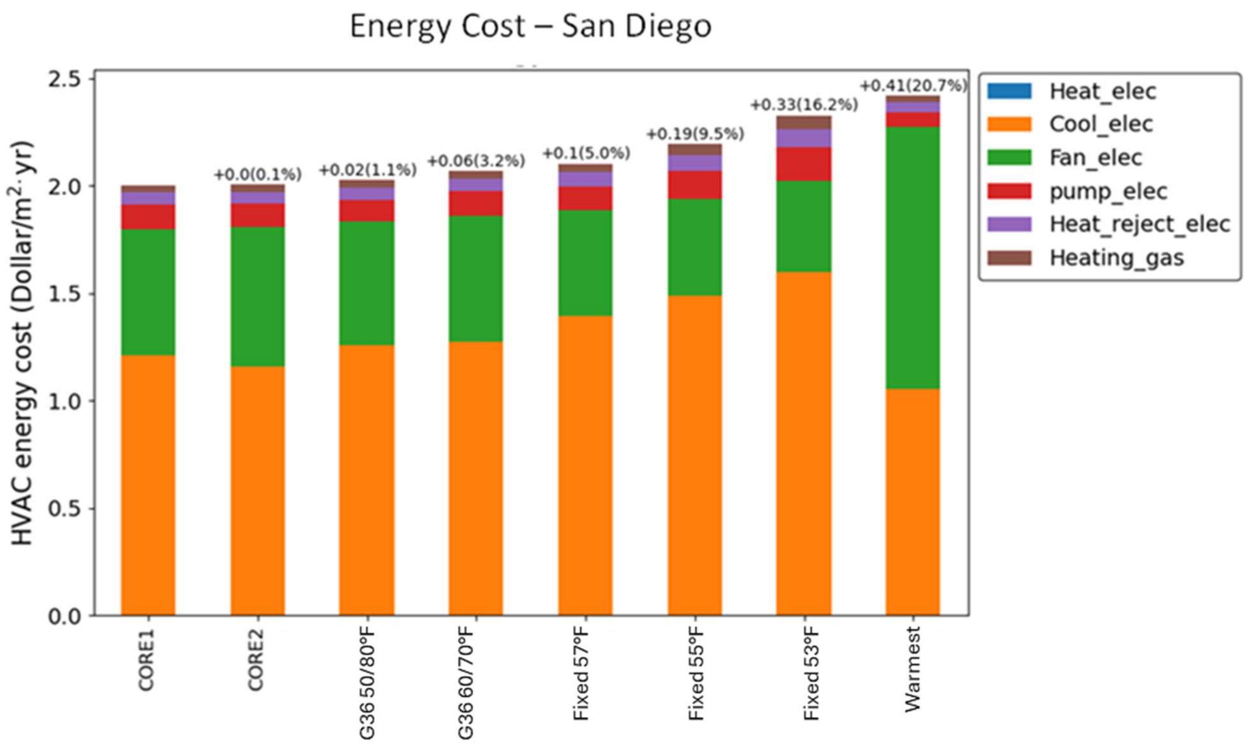

Figure 9: Energy simulation results: Total HVAC energy cost and component breakdown for San Diego climate for CORE and existing SAT control strategies.

Supply Air Temperature, Energy, and Cost Comparison

Figure 10 illustrates how the different strategies—CORE, ASHRAE Guideline 36, and Warmest—control the SAT setpoint at the same OAT. The SATs of all three control strategies converged to the same upper limit of the SAT at the lower range of the OAT, which was less than 10°C or 50°F.

- Warmest used the maximum SAT until the OAT exceeded about 16°C, and its selected SAT was higher than the other two strategies.
- The SAT chosen by ASHRAE Guideline 36 reduced SAT as OAT increased, stopping at the minimum value.
- The SAT set by CORE tended to use lower SATs when the OAT increased in a colder or warmer range, but using higher SATs when the OAT was more moderate, e.g., approximately 12 to 17°C.
- CORE was able to more effectively use free cooling hours to minimize the overall energy cost. Additionally, CORE resulted in more variation in SATs compared to other control methods, suggesting that CORE was dynamically responding to the building operating conditions. The

study found that while OAT was the primary factor, HVAC systems, occupancy, and operational patterns also influenced the best SAT.

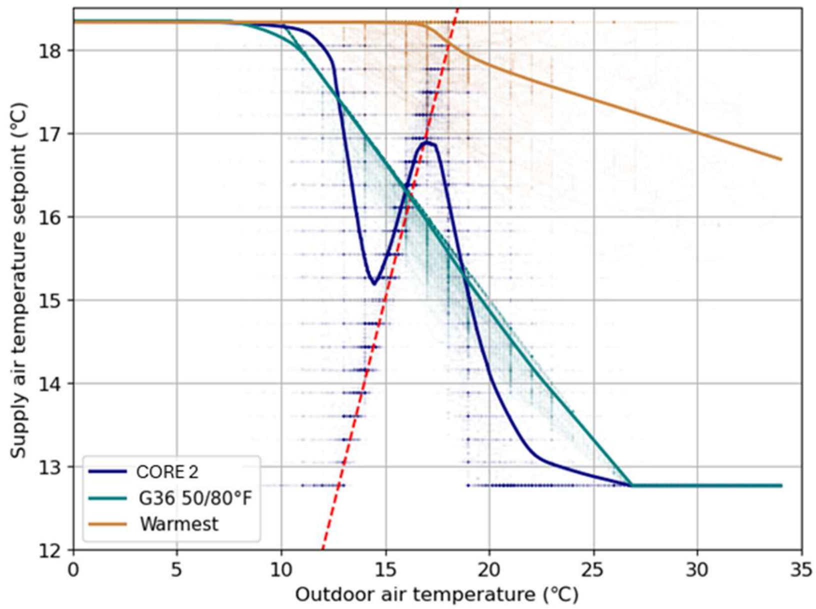

Figure 10: Scatter plot showing the SAT setpoint by OAT for CORE 2, ASHRAE Guideline 36, and Warmest.

Note: CORE 2 appears in blue, ASHRAE Guideline 36 50/80°F OAT limits appear in green, and Warmest appears in orange for the Oakland baseline. Locally estimated scatterplot smoothing (loess) curves were used to depict the overall trend in the data. The red dashed line indicates where SAT equals OAT, i.e., the 100 percent airside economizer line.

The hourly HVAC energy cost, which was normalized by the floor area, was highly correlated with the OAT as expected, as can be seen in the top left pane of Figure 11 below. CORE matched or outperformed ASHRAE Guideline 36 and Warmest across all OAT ranges. The energy savings from CORE were highly dependent on the OAT, with the greatest savings in the range of approximately 16°C to 22°C. In Oakland simulations, the OAT fell within this range for most of the office hours throughout the year. Chilled water, fan, and heating costs are broken down in the remaining subplots of Figure 11 and provide further insights into each control strategy.

- Warmest incurred the highest fan energy costs, the lowest cooling energy costs, and the highest total energy cost, as it led to higher airflow rates due to the selection of higher SATs.
- Both ASHRAE Guideline 36 and CORE opted for lower SATs and lower airflow rates, reducing energy consumption by diminishing fan power, despite a slight increase in cooling energy costs.
- ASHRAE G36 consistently caused higher fan and chilled water costs when the outdoor temperature moved beyond the cold range. However, CORE used higher fan energy use to bring in a larger volume of cooler outdoor air during free cooling hours, a tradeoff that resulted in a lower overall energy cost than ASHRAE Guideline 36. This difference in control behavior corresponded to the distinct SAT selection processes by ASHRAE Guideline 36 and CORE.

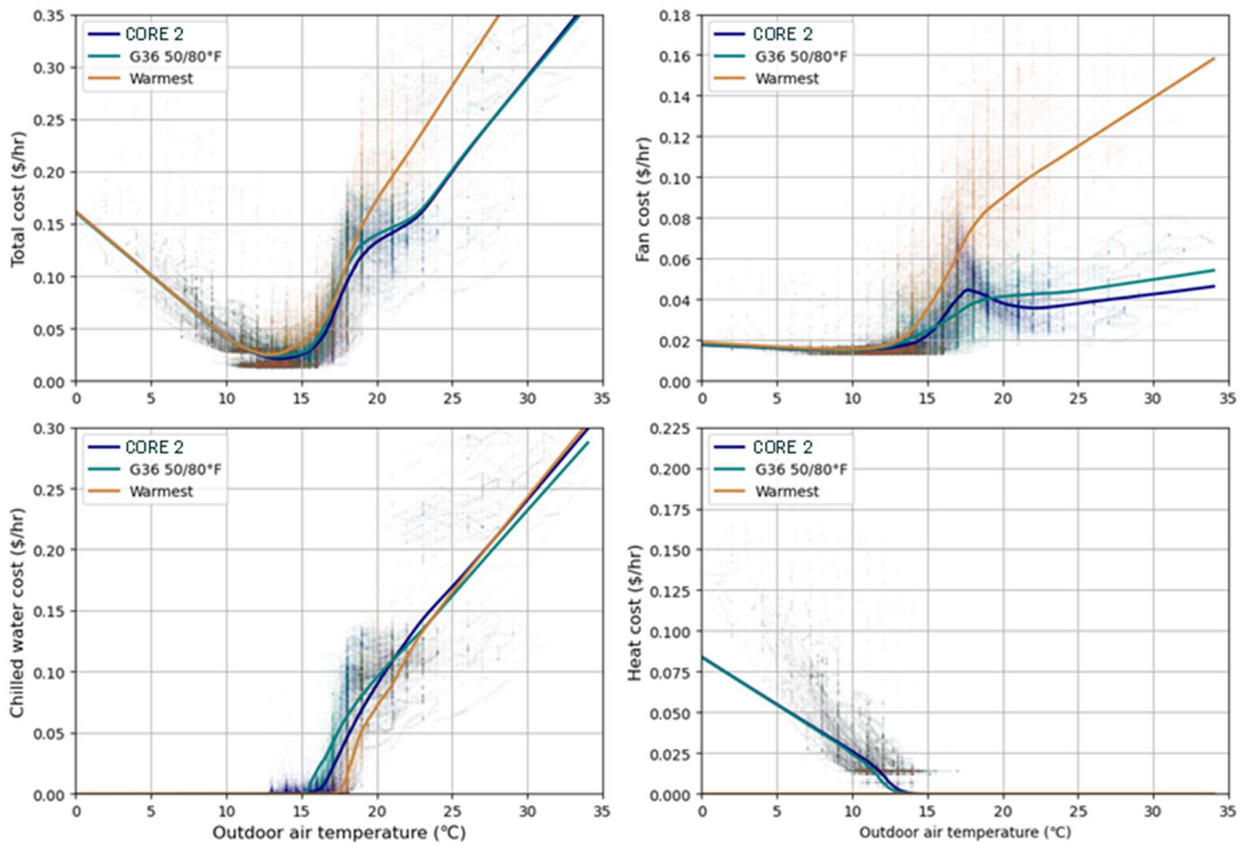

Figure 11: Scatter plot showing the energy cost (per m²) by OAT for CORE 2, ASHRAE Guideline 36, and Warmest.

Note: CORE 2 appears in blue, ASHRAE Guideline 36 appears in green with 50/80°F OAT limits, and Warmest appears in orange for the Oakland baseline scenario. The loess curves were fit to show the overall trend in the data.

Different control strategies selected different SATs throughout a typical cooling day, along with different associated energy costs. The greatest energy savings from CORE occurred when using a lower SAT around 2:00 p.m. to 4:00 p.m., which incurred a very minimal cooling energy penalty but led to substantial fan power savings, as shown in Figure B 1.

Figure 12 illustrates the energy cost comparison of conventional strategies and ASHRAE Guideline 36 to CORE between Oakland, San Diego, and Sacramento. The energy cost savings ($\$/m^2$-year) of CORE and ASHRAE Guideline 36 were very close in Oakland and San Diego, even though the percentage savings were higher in Oakland. This was due to the significantly higher cooling demand in San Diego, where even a smaller percentage reduction in cooling energy use led to substantial cost savings. In contrast, Sacramento saw almost half the dollar savings due to fewer economizing hours.

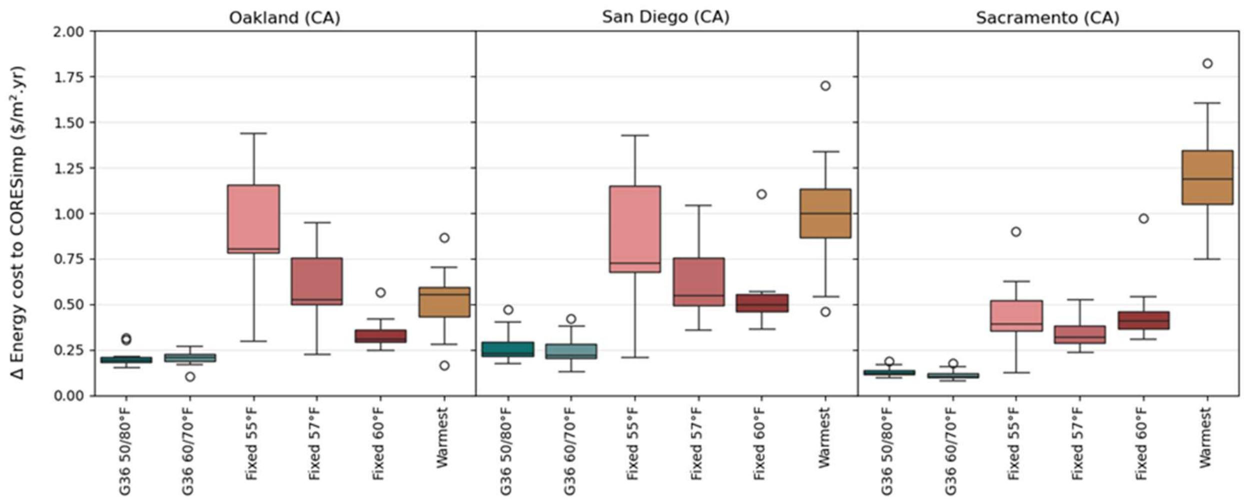

Figure 12: Boxplots comparing energy cost differences of ASHRAE Guideline 36, Fixed SATs, and Warmest vs. CORE 2, aggregating all scenarios for all climate regions.

Note: The boxes depict the 75th and 25th percentiles, with the medians shown as solid lines.

In addition to the total HVAC energy cost, the team also looked at the hourly energy cost and hourly SAT trend data against the OAT for San Diego, as shown in Figure 13. SAT behavior is significantly different between the two strategies. The warmest-zone-based strategy is more expensive at high OATs because it uses warmer supply air. This requires more airflow to meet the zone's cooling needs, which increases the energy cost of running the fans. San Diego has slightly lower total energy cost compared to Oakland, but the overall behavior of SAT and cost in relation to OAT is similar.

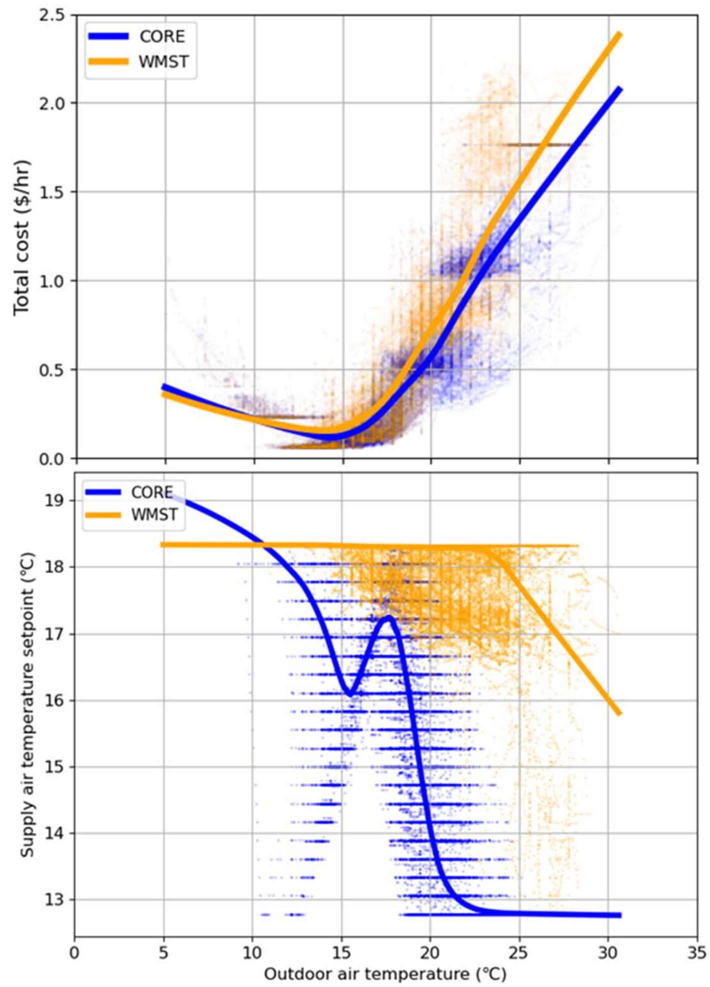

Figure 13: Scatter plots showing total HVAC energy cost and supply air temperature vs. OAT between CORE 1 (CORE) and Warmest (WMST) for San Diego.

Final CORE Algorithm and New Features

Final CORE Algorithm

The project team developed the final CORE algorithm to use in the demonstration sites based on the findings from the previous study and the demonstration site requirements. Appendix A presents the plain-language description of this final CORE sequence of operation used in the three sites. Compared to the initial CORE algorithm, this final version includes a smoothing factor when updating the cooling power. The team introduced this feature to reduce cooling coil oscillation between OFF and the minimum open position during certain cooling load request levels as observed from the previous research study (Raftery, et al. 2018). The team also introduced a cooling power calculation equation for DX cooling systems for demonstration sites with DX cooling, as the initial version of CORE only considered cooling with chilled water.

The final CORE algorithm that the team implemented at the demonstration sites uses zone-level data which corresponds to CORE 1 in Table 1. In addition, the final CORE algorithm presented in Appendix

A features an option to limit the SAT reset range high limit based on outside air dewpoint. This team did not implement this feature in the three demonstration sites as it is only useful in locations and times with high outside relative humidity levels, such as New York in the summer.

The team also reformulated the CORE algorithm into 1) a CORE 2 algorithm, which avoids the need for zone discharge air temperature, and 2) Simple CORE, which avoids using any zone-level data (Table 1). The team did not use these algorithms in the demonstration sites but expects to make them publicly available through a journal publication that is currently under review.

Future Considerations for CORE Algorithm

The team considered the following features for future development:

- Precooling, using the economizer or based upon peak utility rates, could further reduce the peak demand period. During specific high-demand events, having a feature to restrict peak demand may be financially and operationally beneficial.

CORE Implementation Efforts and Issues

The research team and Altura spent considerable effort to implement the CORE algorithm and concluded the CORE algorithm performed as expected.

The research team spent an estimated 184 hours over six months for implementation, including programming support to Altura, implementation review, functional performance testing review, and commissioning, which included issue debugging and resolution support. In addition, Altura spent an estimated 132 hours on CORE programming and commissioning. After completing the initial implementation at three sites, Altura implemented CORE at MOB-4 as a fourth site as a “light touch” implementation with no involvement from the research team. This implementation took only a couple of hours for Altura due to the already available programming for the same supervisory controls structure and due to the prior experience in conducting functional testing and verifying performance. This demonstrates the scalability of CORE implementation in a portfolio of buildings with the same control system structure: The majority of time is required upfront.

The research team and Altura spent an estimated additional 15 hours on site and HVAC controls system issues not directly related to CORE but which impacted its performance. Table 7 summarizes the estimated time consumption by the team under specific implementation tasks for all three buildings.

Table 7: Estimated time spent implementing CORE at the three demonstration sites.

Task	Altura estimated time (hours)	Research team estimated time (hours)
Algorithm and SOO review and support	2	3

Task	Altura estimated time (hours)	Research team estimated time (hours)
Programming	45	68
Program deployment	22	
Commissioning and functional testing	18	113
Testing and debugging	45	
Total	132	184

The team identified several issues with the existing controls and physical equipment at the sites, leading to challenges during the implementation and commissioning phase. These challenges are summarized in the following subsections, including lessons learned for future implementations. The team also provides specific examples and how the team resolved them.

The overall challenges would be present at most buildings and are not specific to the demonstration sites or project partners. This level of effort is not scalable, and as discussed in the Market Scalability Assessment section, this means that CORE should be implemented in a more targeted way.

Algorithm Complexity

There were several issues related to the CORE algorithm programming caused by miscoordination between the implementor and the research team. One reason for this may be the complexity of the CORE algorithm compared to what implementers are accustomed to programming, such as outside air-based SAT control or demand-based SAT control.

Adoption of ASHRAE Guideline 36 has been slow, due in part to complexity and a lack of familiarity among controls contractors and mechanical engineers. ASHRAE Guideline 36 is similar to CORE in its complexity and has been available as an industry best practice, specifically for SAT control, for six years. It also has similar requirements for programming. As a result, CORE may have similar market adoption issues. An example of this issue can be found in Appendix B.

To reduce complexity at individual implementations, BAS manufacturers should consider programming CORE into their programming libraries. Programming libraries are programming logic that BAS manufacturers provide to controls contractors for use in individual buildings. A CORE programming library may alleviate several of the issues listed above, but site implementation is highly dependent on the specific building and the BAS.

To reduce challenges, the team recommends:

- Providing more detailed guidance to the implementer, including an explicit description of the intention of each part of the algorithm.
- Using the simplified algorithm shown in Table 1, which uses fewer data points and reduces the complexity of the implementation.
- Providing an example of “what right looks like” from a programming, testing, and operations perspective. When the implementor and other stakeholders know how the system should be operating, it is easier to verify the effectiveness of CORE throughout implementation and operation.

Existing HVAC System, Building Design, or Control Issues

The team conducted a detailed review of data quality and the general performance of the HVAC system of the buildings before implementation. However, several issues related to the HVAC system and controls and building design still emerged during the commissioning period, which led to CORE not operating as expected. These issues include:

- Non-functional equipment
- Building design that led to high outside airflow
- Issues in the HVAC controller that made it difficult to map and update some variables

An example is with the MOB-1 RTU-1 variable speed compressor. The team observed that RTU-1 did a poor job of meeting the SAT setpoint during occupied times in both CORE and baseline, diminishing CORE’s ability to control SAT and optimize cost, as shown in Figure 14 below. Upon further review, the team found that the DX cooling unit had a non-functioning variable frequency drive for the compressor, leading to the cooling unit going from 0 percent to 25 percent—and then to 50 percent—without intermediate values. This caused SAT values to fluctuate around a setpoint without being able to meet the setpoint, while the cooling unit alternated between two states. The team highlighted this issue to the site project manager, who conducted an on-site investigation; a technician subsequently replaced the compressor. See Appendix B for an additional example.

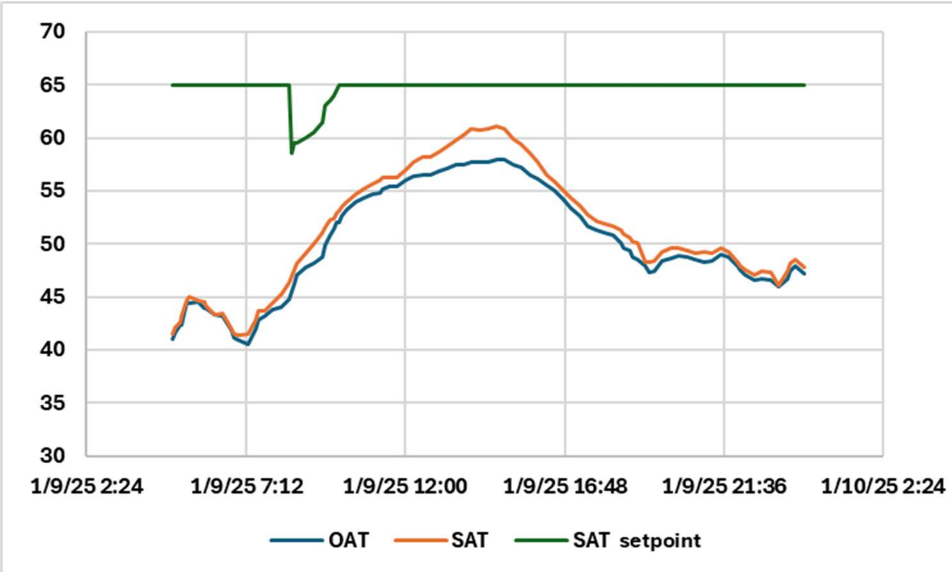

Figure 14: SAT and OAT for MOB-2 RTU 1.

Prior to implementing CORE or other demand-based reset, the team recommends:

- A thorough review of the existing HVAC system, controls, and design
- Performing an existing building commissioning (EBCx)
- Resolving issues identified from the review and EBC

This process itself will yield substantial cost and energy savings in most existing buildings. In terms of system operation, the implementer should evaluate the SAT meeting the setpoint, economizer operation, and a rogue-zone-detecting metric, such as the percentage of hours cooling is requested by each zone. The team also believes that CORE may be best suited as part of a wider control-system upgrade that would scale to resolve these issues more economically, which is further explored in the Market Scalability Assessment section.

Challenges with Working with an Existing Control System

In addition to identifying HVAC system and building issues, the team discovered limitations of the existing control system. Niagara programming follows a wire diagram format that passes on calculated values and logic to next steps. This setup was easy to observe for snapshots of time but it was difficult to investigate how the controller behaved over time, which was necessary for finding the causes of issues. This was partly due to multiple rounds of programming done at the site previously, which complicated the mapping process. In addition, the team uncovered a calculation syntax error that created incorrect CORE calculations and led to CORE not optimizing cost.

The research team performed functional testing of CORE operation after implementation. The intent was to find and correct any operational issues. This testing was useful for identifying several bugs in the programming that were then fixed. However, running these tests was restricted by the capabilities of the system, such as the ability to override variables due to equipment restrictions, which could make it difficult to identify issues. The team also found several issues by performing manual calculations of each CORE variable using downloaded measured trend data. It may be useful to do the manual calculations early on as part of commissioning. Performing a detailed review through live access to programming may help identify issues faster. Using a simpler CORE algorithm may also help, as this would require less programming and use a smaller number of variables.

An example of this issue came up for the cooling power (Pc) syntax. Pc values were about 10 times larger than what we expected when looking at equipment capacity and load. This led to calculated cooling power and cost being consistently much higher than heating and fan power and cost, causing the CORE algorithm to push the SAT setpoint to higher values to reduce cooling power and cost. To find more details about the Pc syntax and an additional example, see Appendix B.

The team recommends:

- A thorough review of the control system functionality, including accessibility to software for programming and configuring the control system.
- Speaking to facility controls technicians about existing issues can help identify problems beforehand, along with coordinating with the control system service contractor on what is needed to implement.
- Developing a standard pre-check of system requirements, coupled with functional testing and commissioning procedures to verify CORE performance prior to handoff to facility operations.

Energy and Cost Savings Results

Energy and Cost Savings Summary

Overall, all three buildings showed energy and cost savings using CORE. Table 8 below summarizes the measured energy and cost savings for MOB-1, MOB-2, and MOB-3. Note that the team calculated energy cost using only the per energy unit rate and we did not add fixed or demand charges.

MOB-1: CORE reduces energy consumption by 5 to 16 percent and cost by 1 to 18 percent based on the RTU level heating, cooling, and fan (HCF) data.

MOB-2: CORE reduces gas energy consumption by 13 percent at the whole-building level, but with poor statistical evidence. HCF energy, whole-building-level electricity, and total energy are similar between Baseline and CORE.

MOB-3: CORE reduces energy consumption by 15 percent and reduces cost by 20 percent, according to heating, cooling, and fan data.

The magnitude of HVAC cost, especially considering the floor area normalized annual cost which comes out to be between $ 0.05 to 0.16 per square feet per year, is low compared to typical buildings in California where values around $ 0.5 per square feet per year would represent a well performing building. The reason for this is the low electricity TOU rates applicable to these buildings. In a typical commercial building in California, these rates are higher, and CORE would have higher absolute energy cost savings.

Table 8: Energy and cost savings summary for MOB-1, MOB-2, and MOB-3.

	Measured Mean Daily HVAC Energy				Measured Mean Daily HVAC Cost			
Energy type	Baseline (kWh)	CORE (kWh)	Savings	P-value	Baseline ($)	CORE ($)	Savings	P-value
MOB-1 HCF RTU1	167	148	11%	0.058	4.15	4.10	1%	0.469
MOB-1 HCF RTU2	97	92	5%	0.288	2.96	2.63	11%	0.265
MOB-1 HCF RTU3	227	190	16%	0.042	6.57	5.40	18%	0.156
MOB-2 Total HCF	481	470	2%	0.410	14.96	13.39	10%	0.154
MOB-2 Building Elec+Gas	2307	2142	7%	0.163				
MOB-2 Building Elec	1337	1353	-1%	0.558				
MOB-2 Building Gas	942	822	13%	0.126				
MOB-3 Total HCF	229	178	22%	0.021	6.63	5.33	20%	0.065
MOB-3 HVAC Elec+Gas	497	422	15%	0.059				
MOB-3 HVAC Elec	441	363	18%	0.074				
MOB-3 HVAC Gas	72	71	1%	0.487				

MOB-1

SAT CHARACTERIZATION

The research team evaluated CORE operation and cost of heating, cooling, and fan operation using trend data for each RTU/AHU at each demonstration site. A representative day for CORE operation that indicates how the SAT setpoint is driven based on the calculated energy cost is shown in Figure 15. For most of the day, cost for operating at the current SAT+0.5°F has the lowest cost, and this leads to the algorithm resetting SAT setpoint to increase until it hits the limit at 65°F. The team did not evaluate above 65°F but doing so may allow CORE to run at even lower costs during the times when the SAT setpoint hits 65°F.

During midday, when outside air is warmer, reducing SAT becomes cheaper (current SAT – 0.5°F). This is likely because with warmer outside air coming into the RTU, cooler supply air at a lower flow rate leads to lower total cost. Around late afternoon, when the outside air becomes cooler, using the cool outside air to reduce the mechanical cooling requirement becomes more cost effective, so the SAT setpoint is increased. Throughout the day, total cooling requests remain below the number of ignores, meaning that CORE always uses cost optimization path.

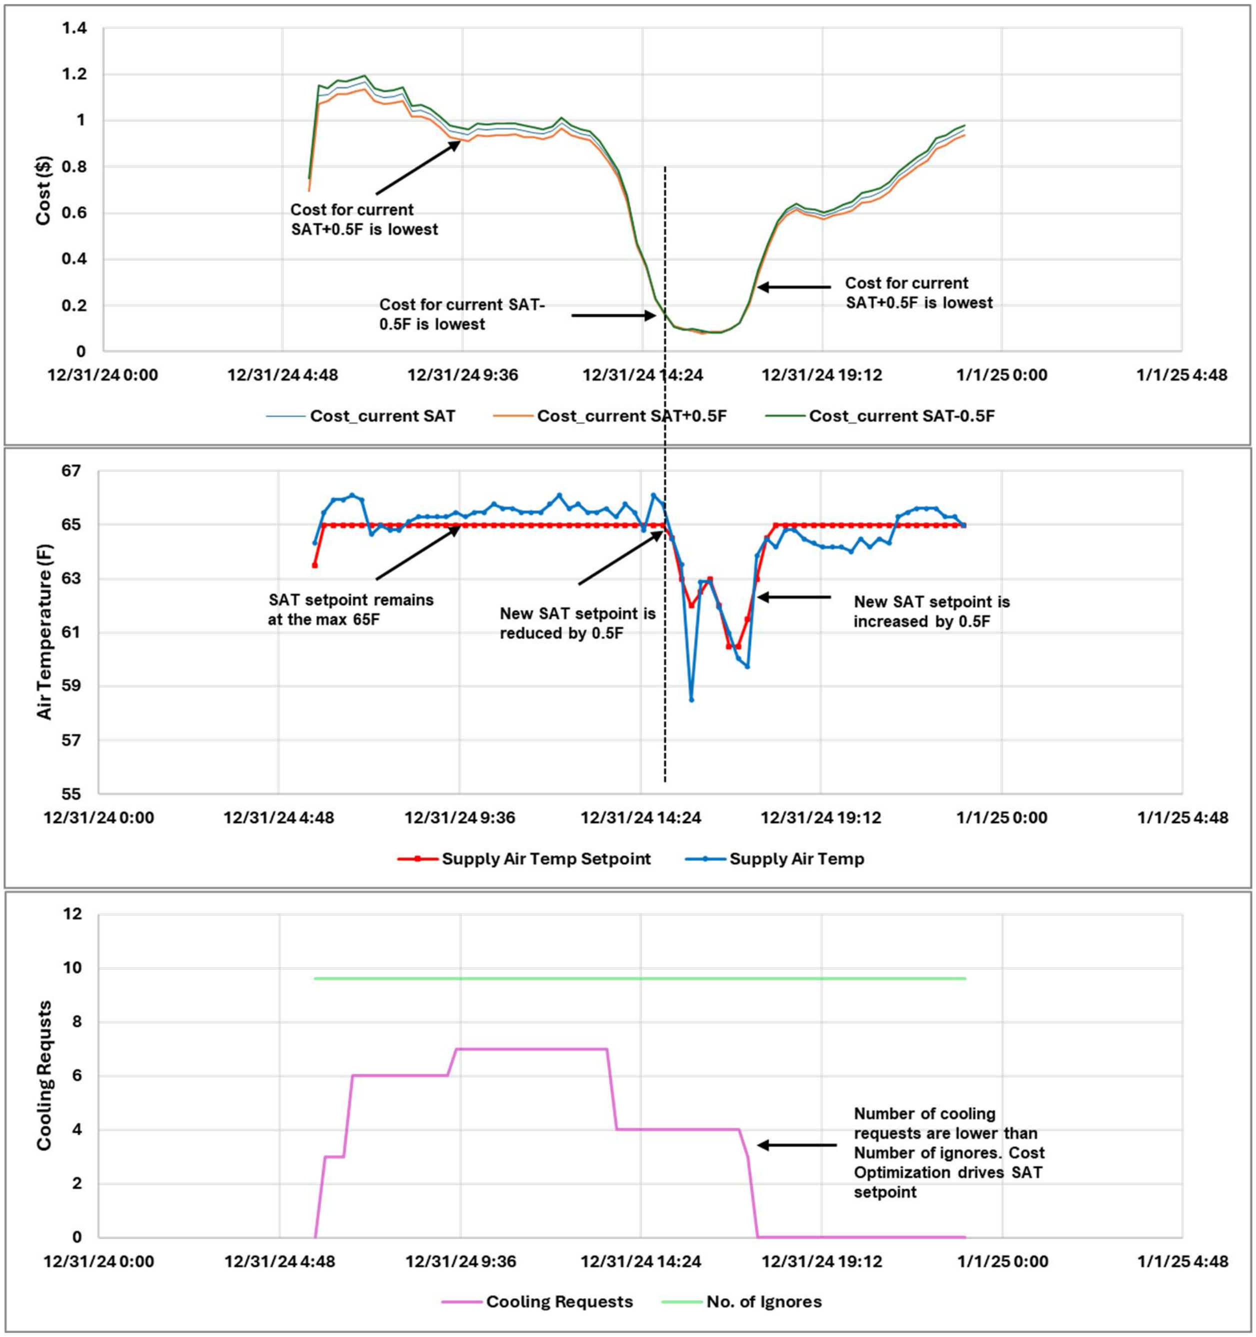

Figure 15: Annotated trend data for CORE operation in MOB—1RTU-3 showing cost optimization when cooling results are less than the number of ignores.

Figure 16 shows the cost data and SAT data for a day with CORE operation where during midday, zones requested cooling. This can be seen from the total cooling requests exceeding the ignores limit, leading to a sharp decrease in the SAT setpoint (and SAT) to meet demand. Despite the CORE algorithm showing cost for current SAT+0.5°F being the lowest, the SAT was reduced to meet cooling needs. This suggests increasing SAT to optimize cost. To meet the cooling demand, the CORE algorithm uses T&R logic similar to ASHRAE Guideline 36.

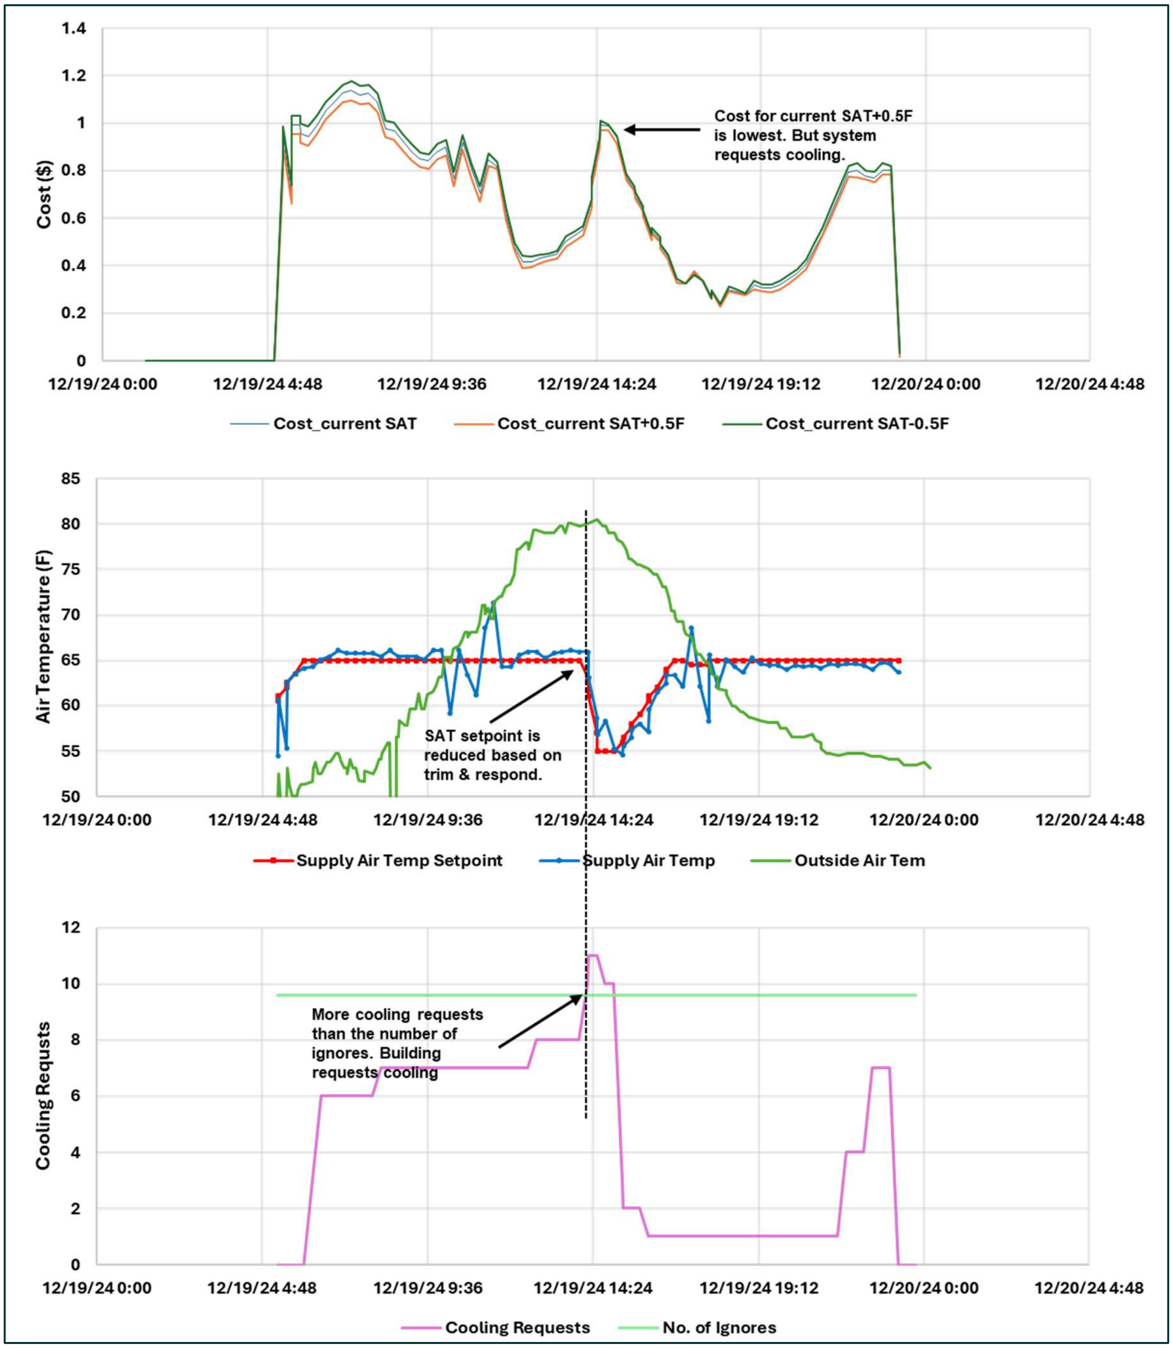

Figure 16: Annotated trend data for CORE operation in MOB-1—RTU-3 showing cooling requests driving supply air temperature down.

The team compared SAT setpoint behavior on a baseline day to a CORE-enabled day. Figure 17 shows two days in January 2025 that had similar OAT conditions. The SAT setpoint behavior is highlighted and shows that CORE optimizes for cost using the heating, cooling, and fan energy cost. The decrease in SAT setpoint shown by the dashed line indicates that the CORE algorithm optimized the total cost by increasing cooling cost and decreasing fan cost In comparison, the baseline day controlled SAT based on the cooling requests, with the number of ignores at two, and the total HVAC cost increased to meet the cooling demand.

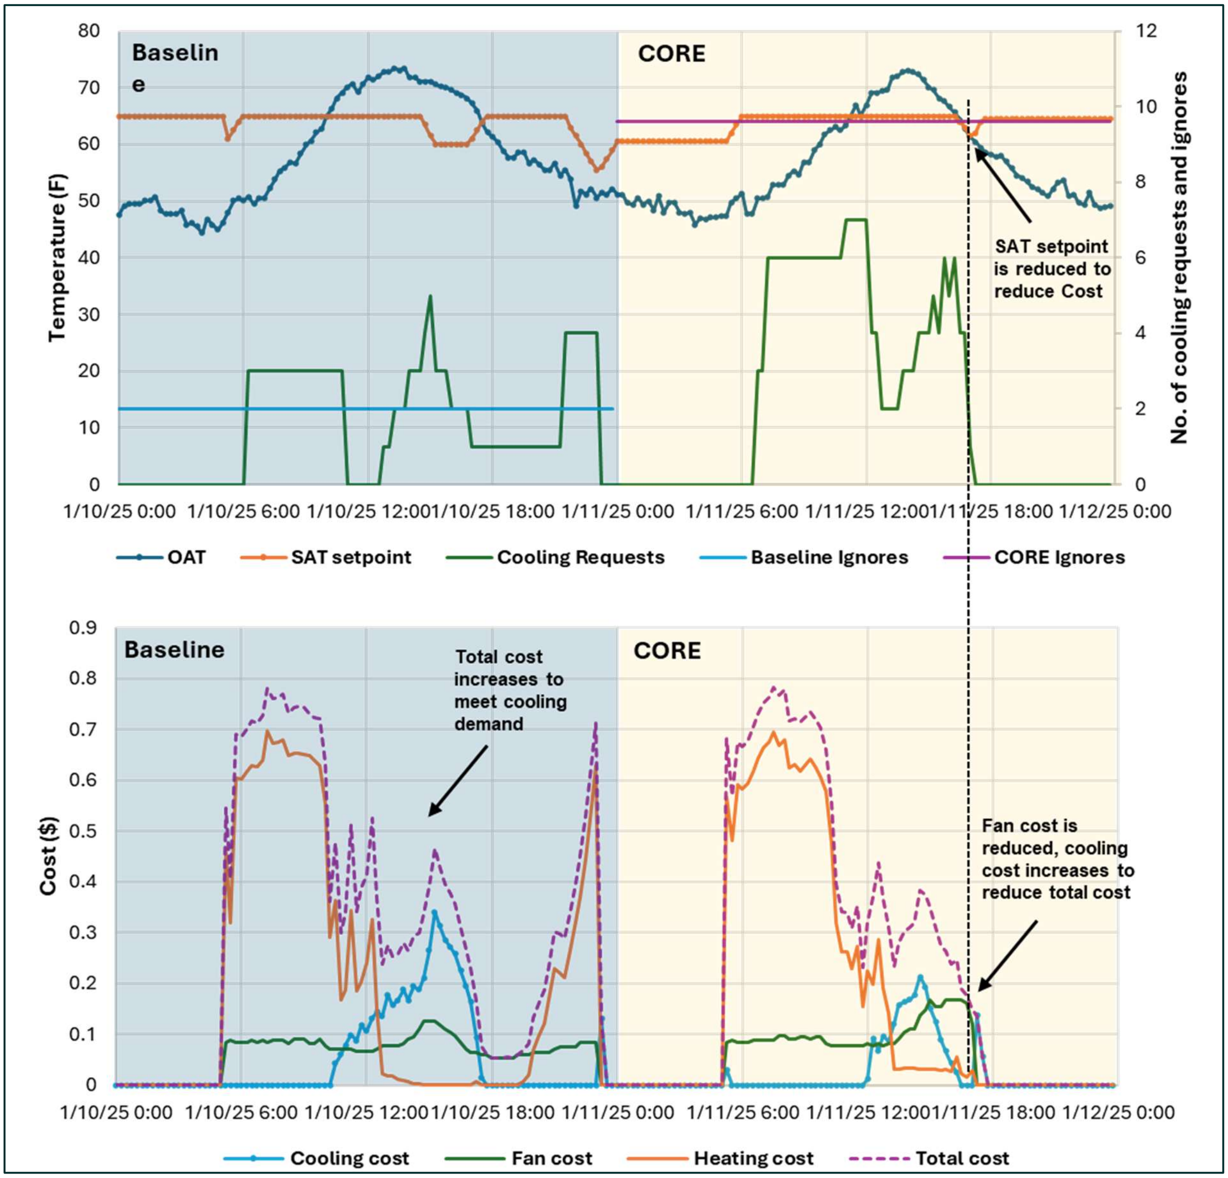

Figure 17: Comparison of Baseline and CORE operation for cost vs. supply air temperature setpoint.

Figure 18 shows the SAT setpoint against outside air for the same two days: January 10 and 11, 2025. The baseline strategy tended to set cooler SAT setpoints at warmer OATs, while CORE achieved the thermal demand by balancing cooling and supply airflow rate.

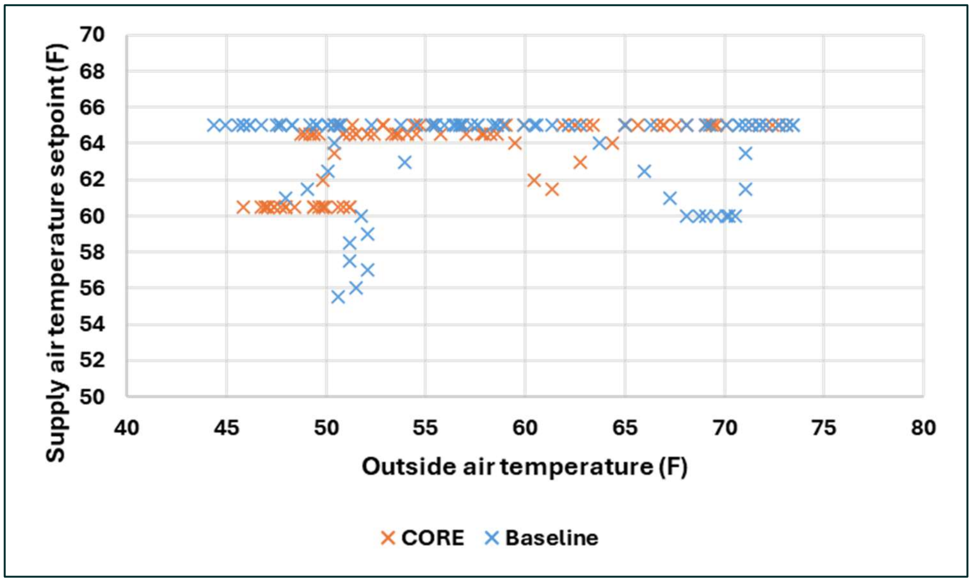

Figure 18: Comparison of supply air temperature setpoint behavior with OAT for Baseline vs. CORE—MOB-1.

Figure 19, Figure 20, and Figure 21 show the histogram of SAT setpoint distribution using 5-minute data for baseline and CORE for each RTU throughout the monitoring period. RTU 1 and 3 have clear differences between SAT setpoint for Baseline and CORE. This aligns with the energy savings of 11 percent and 16 percent, which are described in the Energy and Cost Summary section. RTU 2, which has a similar SAT setpoint distribution for both Baseline and CORE, has smaller energy savings of 5 percent.

RTU 1 runs 24/7 to provide cooling to a few zones with medical equipment, but during the nighttime, the majority of the zones do not have any heating demand. Therefore, CORE frequently runs at the lowest SAT setpoint and saves both fan energy and reheat energy without compromising comfort.

RTU 3 runs at warmer SAT setpoints under CORE to be able to meet comfort requirements during occupied times and balances this with higher air flow rates. The total energy and cost is less in CORE compared to Baseline. RTU 2 has similar SAT setpoint distributions for both Baseline and CORE and this aligns with the much closer average daily energy consumption between strategies for RTU 2 as seen below.

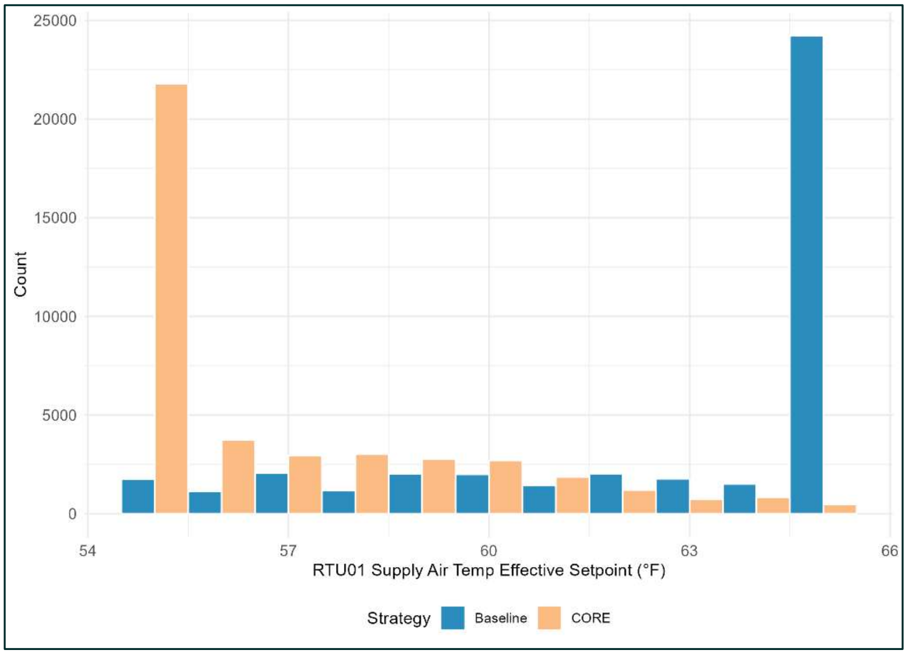

Figure 19: Supply air temperature setpoint distribution for Baseline and CORE—RTU 1—MOB-1.

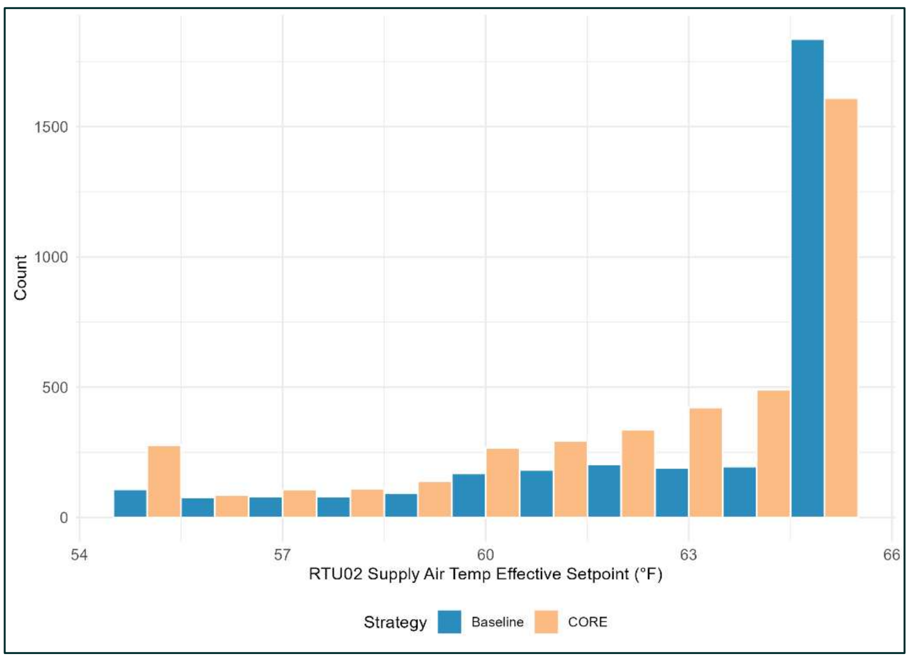

Figure 20: Supply air temperature setpoint distribution for Baseline and CORE—RTU 2—MOB-1.

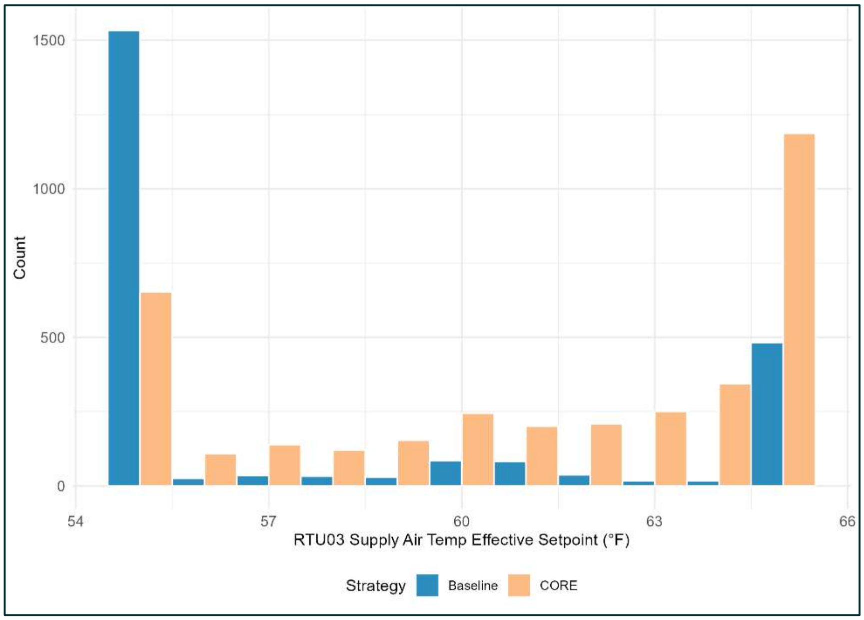

Figure 21: Supply air temperature setpoint distribution for Baseline and CORE—RTU 3—MOB-1.

ENERGY AND COST CHARACTERIZATION

Due to the issues related to existing controls discussed earlier, MOB-1 had several periods where CORE was not running correctly for each RTU. Based on the data processing plan, the team removed these periods from the analysis. The results presented for MOB-1 focus on energy and cost for each of the three RTUs separately since the team does not think total HVAC energy values are accurate for a meaningful period.

Figure 22, Figure 23, and Figure 24 below show the total heating, cooling, and fan energy consumption by each RTU against the OAT. For all three RTUs, CORE uses less energy for moderate OATs. For RTU 1, this is 55 to 62ºF; for RTU 2, 60 to 75ºF; and for RTU 3, below 70ºF. CORE uses higher SAT setpoints and balances the air flow rate, or fan energy, to achieve the lowest cost.

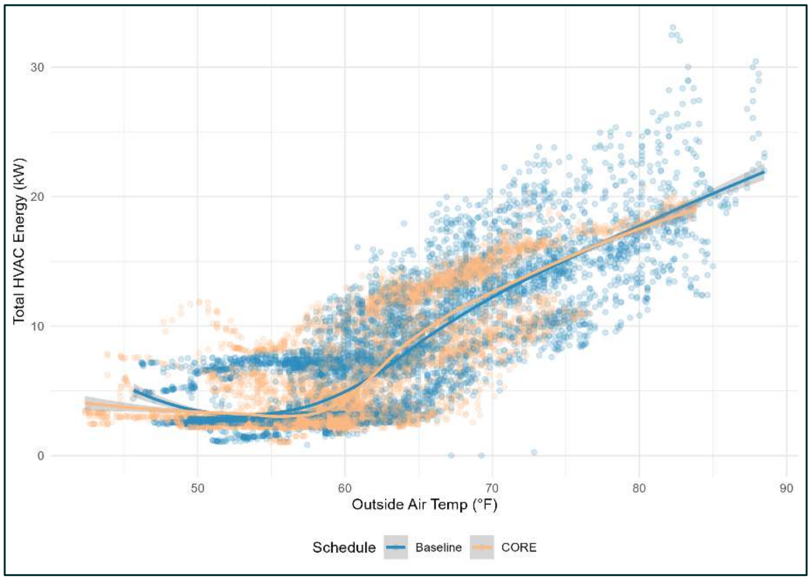

Figure 22: Total heating, cooling, and fan energy consumption vs. OAT for RTU 1—MOB-1.

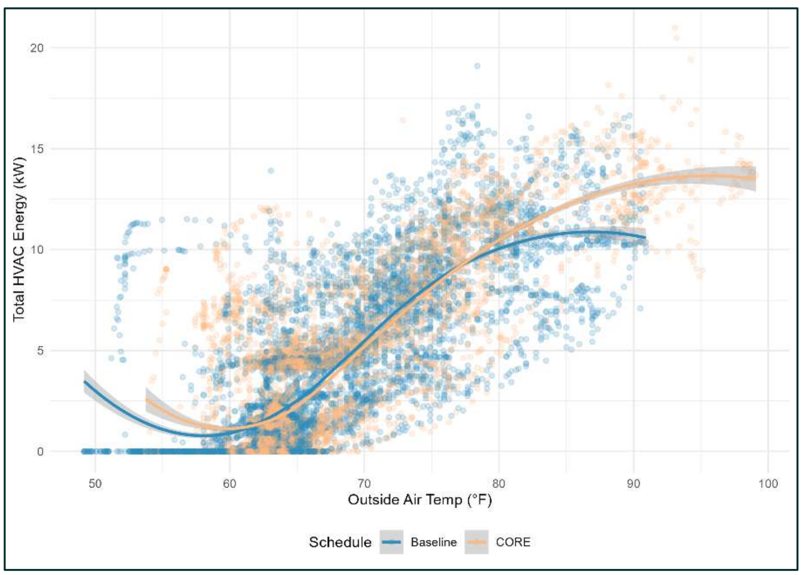

Figure 23: Total heating, cooling, and fan energy consumption vs. OAT for RTU 2—MOB-1.

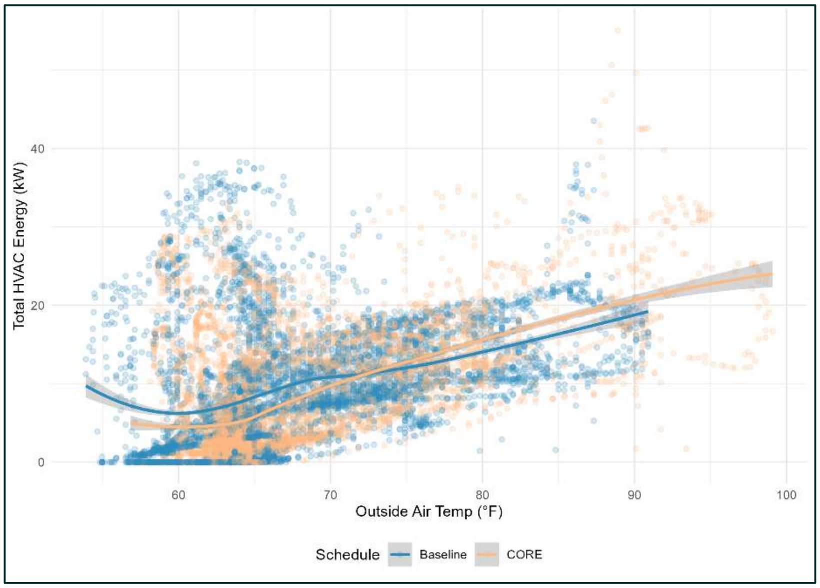

Figure 24: Total heating, cooling, and fan energy consumption vs. OAT for RTU 3—MOB-1.

Similarly, Figure 25, Figure 26, and Figure 27 show the total heating, cooling, and fan energy cost for the RTUs against the OAT.

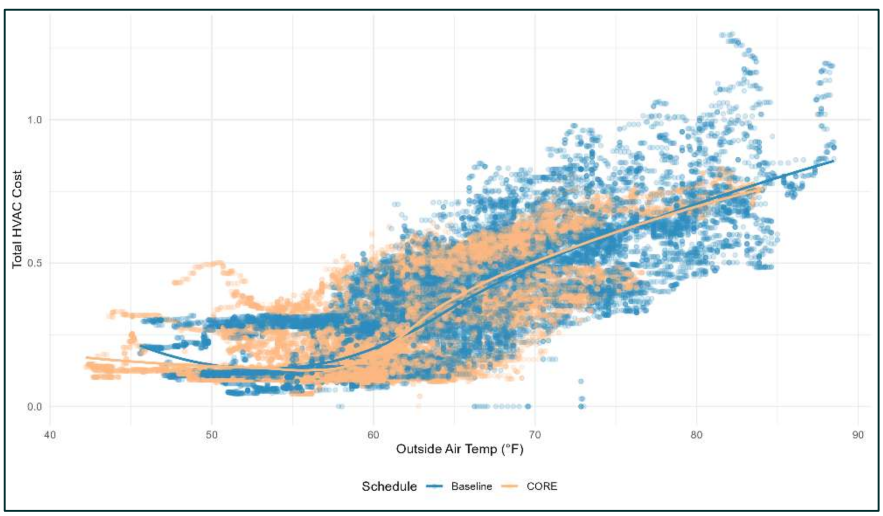

Figure 25: Total heating, cooling, and fan energy cost vs. OAT for RTU 1—MOB-1.

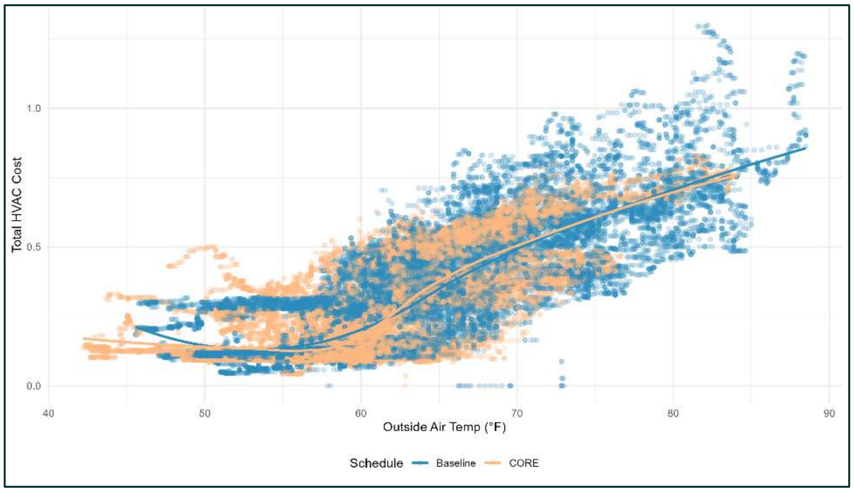

Figure 26: Total heating, cooling, and fan energy cost vs. OAT for RTU 2—MOB-1.

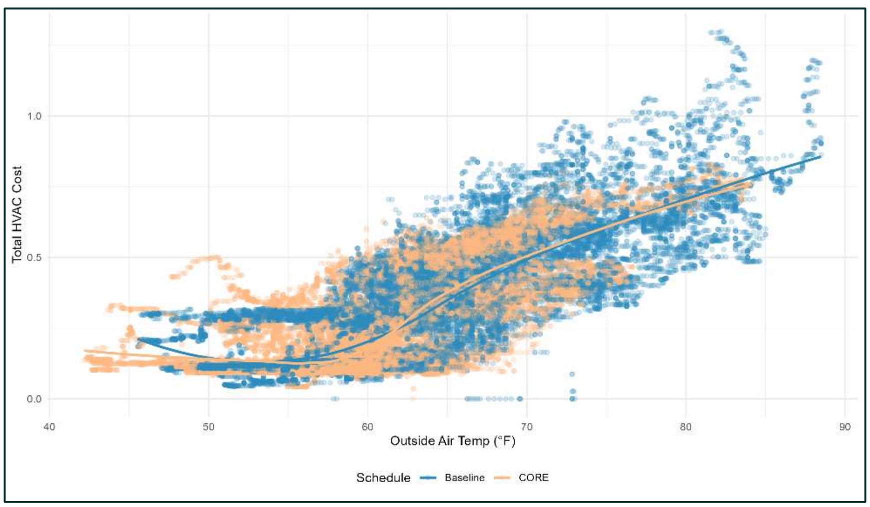

Figure 27: Total heating, cooling, and fan energy cost vs. OAT for RTU 3—MOB-1.

ACTUAL ENERGY AND COST DIFFERENCE

Table 9 shows the distribution of daily energy use for Baseline and CORE measurements for each RTU. Figure 28 illustrates the distribution of daily energy consumption through a violin plot for RTU 1. The analysis includes a p-value that represents the probability of the null hypothesis being true. In this context, the null hypothesis states that the mean daily energy usage from the CORE measurement set is greater than or equal to that of the Baseline measurement set.

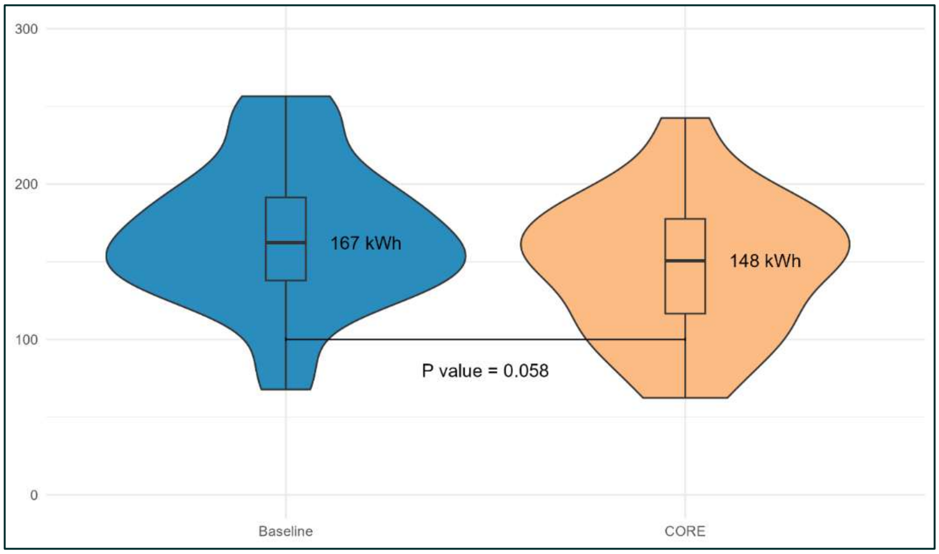

Figure 28: Violin plot comparing daily energy consumption distributions for Baseline and CORE RTU 1—MOB-1.

Note: For additional distribution plots, see Figure B 2, Figure B 3, Figure B 4, Figure B 5, Figure B 6, and Figure B 7

Table 9: Daily energy use for Baseline and CORE measurements for each RTU.

RTU	Baseline (kWh)	CORE (kWh)	P-value
RTU 1 – MOB-1	167	148	0.058
RTU 2 – MOB-1	97	92	0.288
RTU 3 – MOB-1	227	190	0.042

Low p-values for RTU 1 and 3 (less than 0.1) indicate strong evidence that CORE consumes less energy than Baseline. For RTU 2, the difference of the mean daily energy consumption between CORE and Baseline is low, and the p-value is high, indicating weak evidence to support the claim that CORE used less energy. This aligns with Figure 20 above that shows the SAT setpoint distribution is similar for both strategies in RTU2.

Table 10 shows a similar comparison of HVAC cost for the three RTUs. P-values are high in all three cases, indicating weak evidence to support the claim CORE saved energy cost. However, the mean daily cost is lower for CORE in all three RTUs. MOB-1 is an all-electric building, however, there is a difference between energy consumption and cost results due to the energy use at different times of the day where the TOU utility rate is different.

Table 10: Comparison of daily HVAC cost for the three RTUs

RTU	Baseline	CORE	P-value
RTU 1 – MOB-1	$4.15	$4.10	0.469
RTU 2 – MOB-1	$2.96	$2.63	0.265
RTU 3 – MOB-1	$6.57	$5.10	0.156

NORMALIZED ENERGY USE IMPACT

The energy consumption models are based on TOWT models, as discussed in the Energy Savings Analysis section. The team used the calculated energy consumption for heating, cooling, and fan energy, and the measured total based on power meter data, together with the OAT at each RTU and AHU.

Table 11 presents the TOWT energy model statistics for RTUs and the total energy for each site. The HCF energy indicates the heating, cooling, and fan energy calculated and used by the CORE algorithm. Total HVAC energy is the measured HVAC energy based on metered power. Some of the models do not meet all of the criteria for acceptable model fit highlighted in theEnergy Savings Analysis☐ section where acceptable values are in green. The team believes this is caused by the limited data used, as well as the issues in CORE implementation discussed earlier.

Table 11: Model fit statistics for selected RTUs and buildings using TOWT method

	Baseline				CORE			
Energy type	No. of days	R2	CVRMSE%	NMBE %	No. of days	R2	CVRMSE %	NMBE %
MOB-1 HCF RTU1	27	0.92	25.58	1.29E-12	28	0.93	22.78	7.0E-13
MOB-1 HCF RTU2	20	0.79	22.6	-4.0E-12	19	0.80	23.71	-3.0E-12
MOB-1 HCF RTU3	22	0.84	40.55	3.3E-12	22	0.72	64.04	-7.8E-11
MOB-2 RTUs Total HCF	39	0.78	20.48	-5.0E-12	40	0.74	25.51	7.1E-12
MOB-2 Building Elec+Gas	39	0.65	1.07	-1.5E-12	40	0.67	12.44	7.0E-13

	Baseline				CORE			
MOB-2 Building Elec	39	0.92	10.26	2.2E-13	40	0.91	11.51	2.2E-14
MOB-2 Building Gas	39	0.77	24.57	7.9E-12	40	0.78	28.9	-2.0E-12
MOB-3 AHUs Total HCF	48	0.32	68.25	-4.3E-13	49	0.33	68.36	5.3E-12
MOB-3 HVAC Elec+ Gas	48	0.72	39.21	1.2E-11	49	0.68	47.75	-6.1E-12
MOB-3 HVAC Elec	48	0.83	35.42	5.8E-12	49	0.81	43.15	-5.3E-12
MOB-3 HVAC Gas	48	0.56	219	3.9E-11	49	0.50	225	1.3E-01

Figure 29, Figure 30, and Figure 31 show the actual energy consumption versus TOWT model prediction for each RTU. The trends generally indicate the expected pattern where energy consumption is higher at higher OATs. However, the data available for MOB-1 is limited and thus the team could not evaluate the normalized annual energy consumption based on the model due to the limited input conditions (OAT).

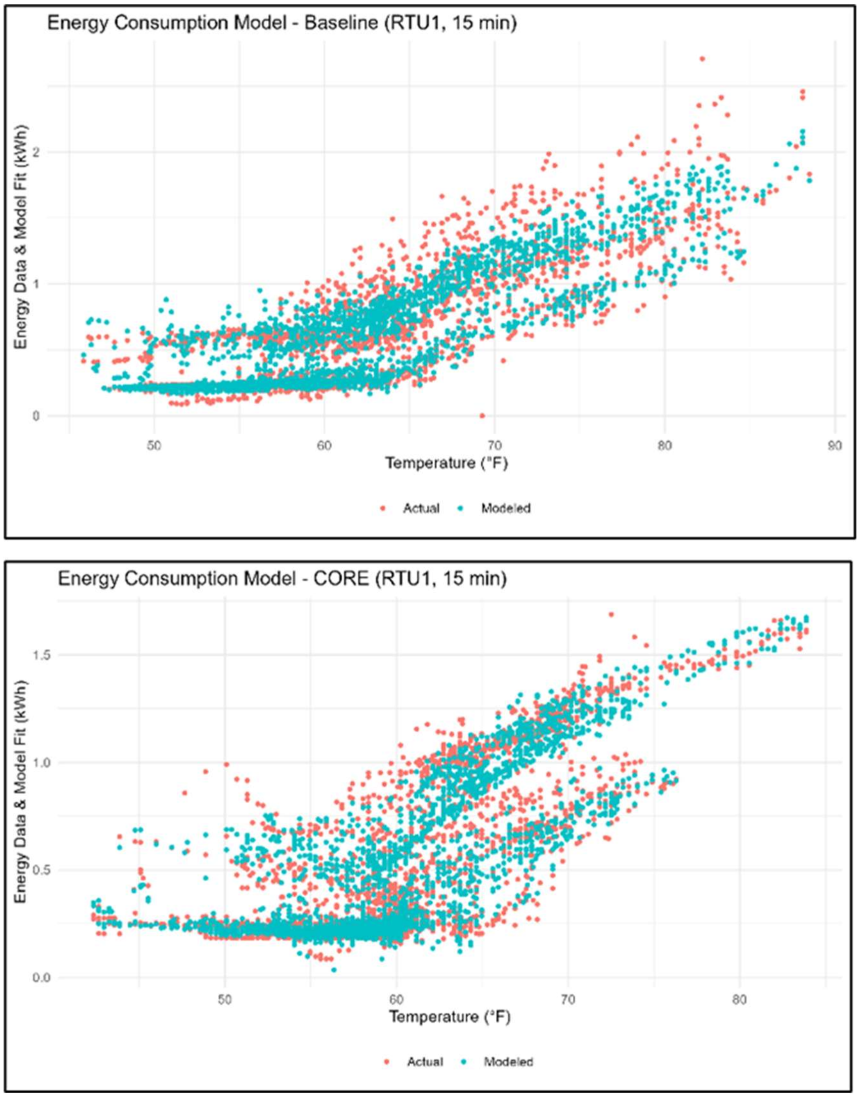

Figure 29: Measured energy vs. model-predicted energy for Baseline (top) and CORE (bottom) conditions—RTU 1—MOB-1.

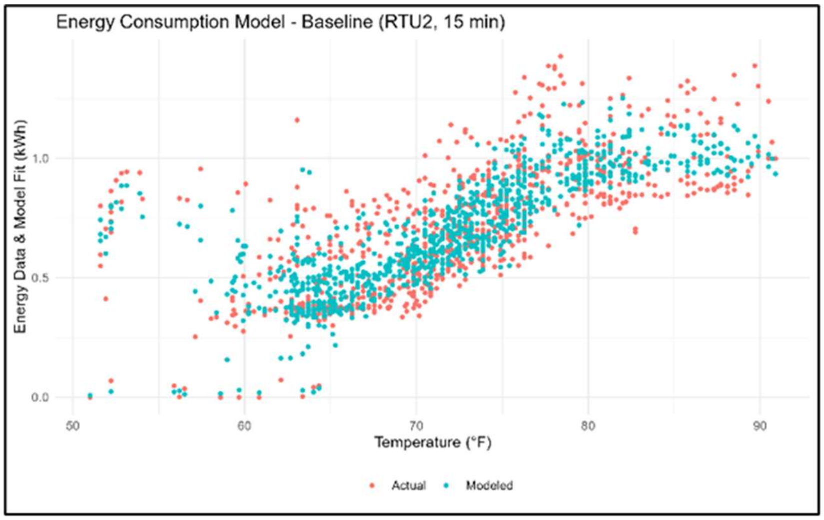

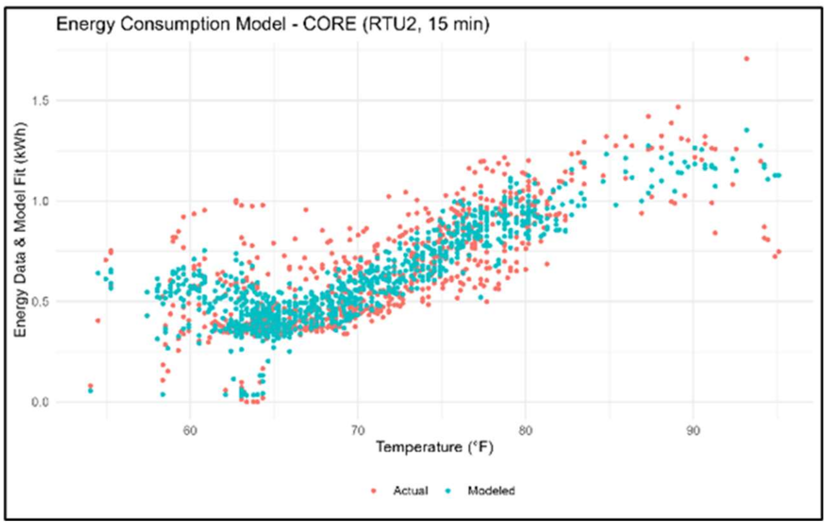

Figure 30: Measured energy vs. model-predicted energy for Baseline (top) and CORE (bottom) conditions—RTU 2—MOB-1.

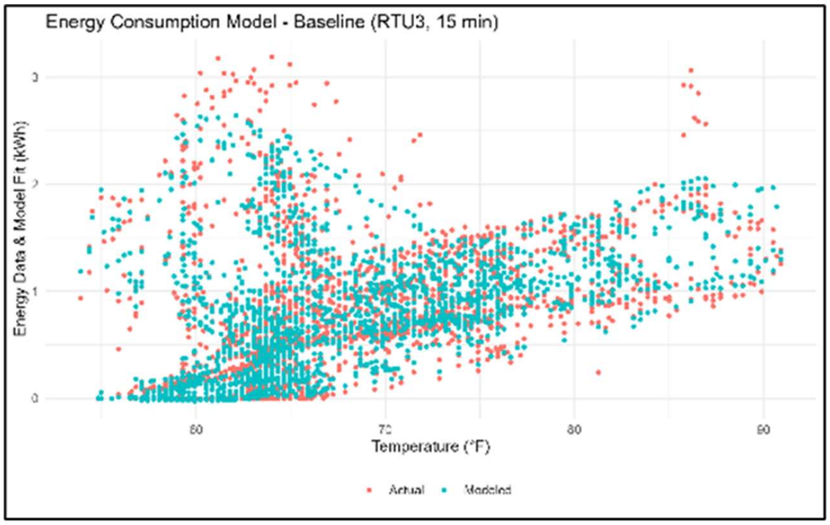

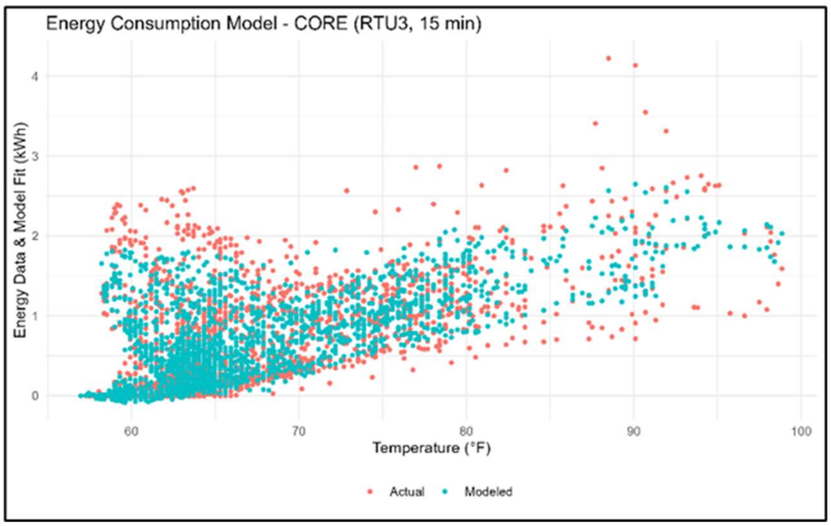

Figure 31: Measured energy vs. model-predicted energy for Baseline (top) and CORE (bottom) conditions—RTU 3—MOB-1.

MOB-2

MOB-2 consists of three RTUs; however, RTU 2 and 3 are combined ("RTU 2/3") to supply air to the same plenum. Therefore, in the analysis below, they are considered as a single system.

SAT CHARACTERIZATION

Figure 32 and Figure 33 show the histogram of SAT setpoint for Baseline and CORE for the two RTUs at MOB-2. Under Baseline, MOB-2 limits the SAT reset range between 55 to 58°F, whereas CORE resets it between 55 to 65°F. This allows CORE to operate at higher SAT setpoints reducing cooling energy requirement.

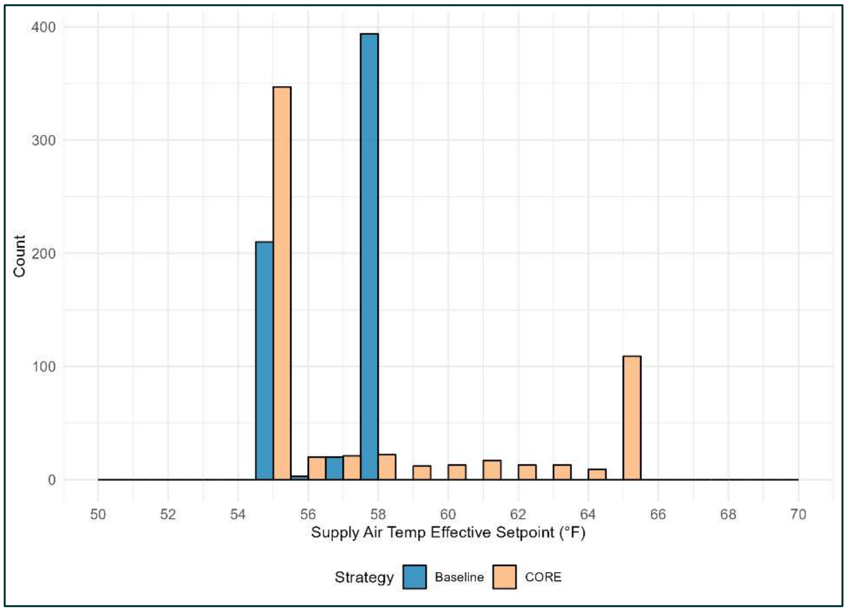

Figure 32: Supply air temperature setpoint distribution for Baseline and CORE—RTU 1—MOB-2.

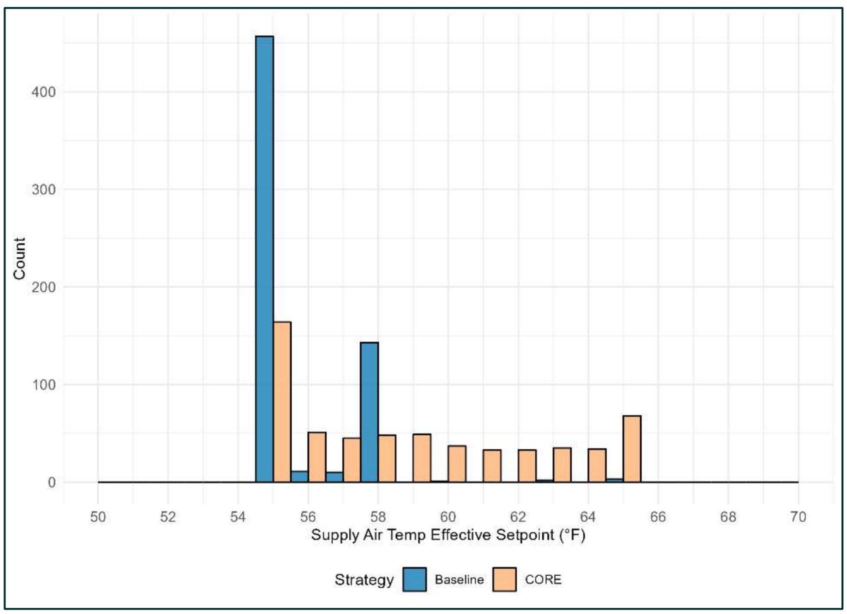

Figure 33: Supply air temperature setpoint distribution for Baseline and CORE—RTU 2/3—MOB-2.

Figure 34 and Figure 35 show the time -of-day distribution of the SAT setpoint for the two RTUs at Baseline and CORE.. The trend lines suggest that in RTU 1warmer setpoints are used in the afternoon when the cooling load is likely higher, allowing for reduced cooling energy, which is partially offset by increased fan energy to meet comfort requirements. For RTU 2/3, CORE always sets higher SAT setpoints, thus reducing cooling energy consumption.

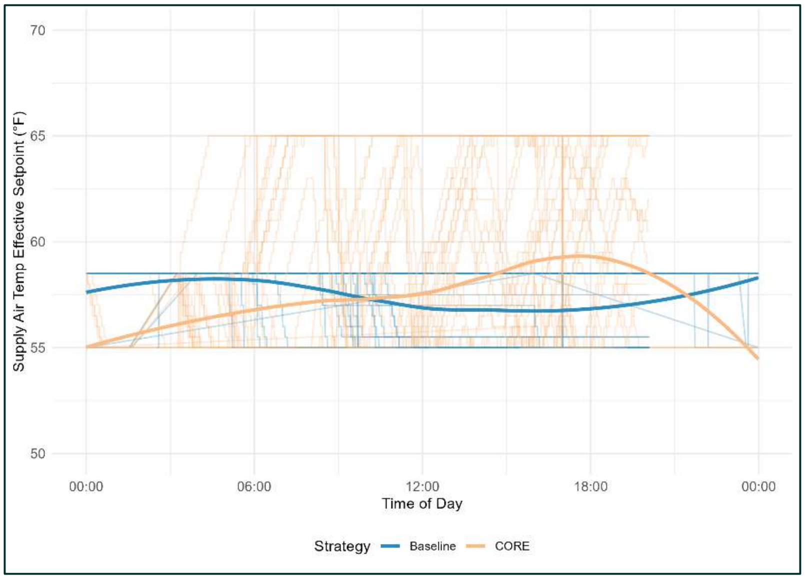

Figure 34: Supply air temperature distribution by time-of-day for RTU 1—MOB-2.

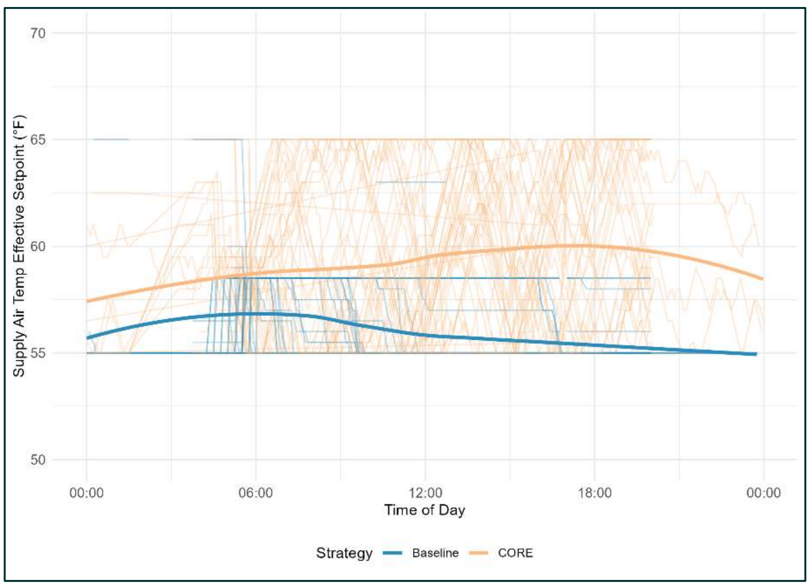

Figure 35: Supply air temperature distribution by time-of-day for RTU 2/3—MOB-2.

ENERGY AND COST CHARACTERIZATION

Figure 36 shows the heating, cooling, and fan power consumption for RTU 1 and RTU 2/3 against the OAT. As expected, there is higher energy use at higher and lower OATs due to more cooling and heating requirements.

Energy consumption for Baseline and CORE is similarly distributed, with CORE spending less energy at higher temperatures.

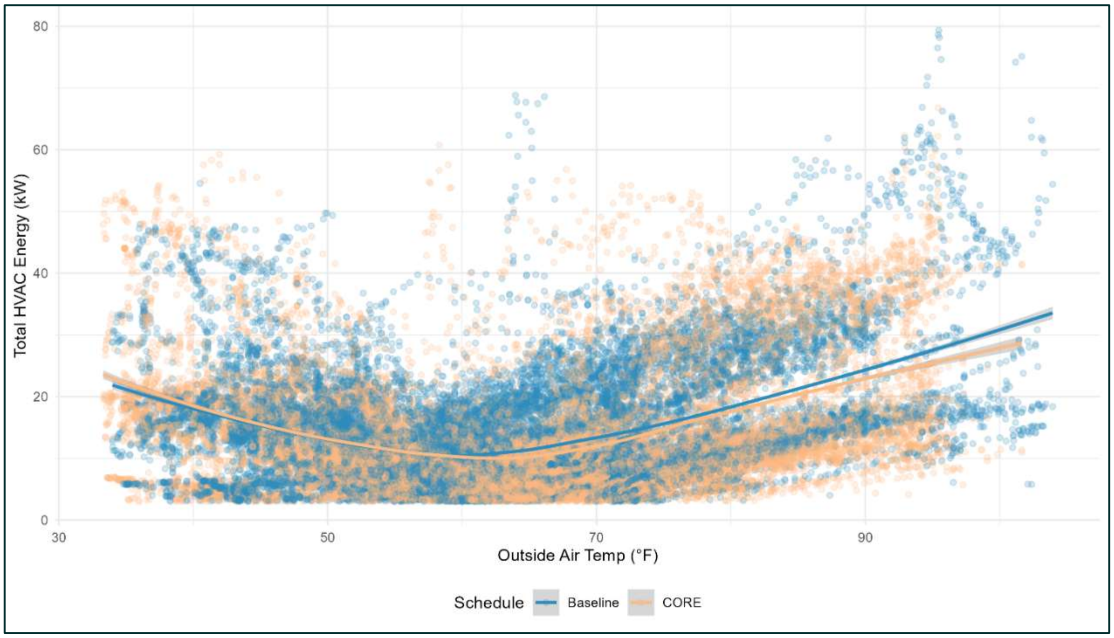

Figure 36: HVAC heating, cooling, and fan energy consumption vs. OAT for combined RTU 1 and RTU 2/3—MOB-2.

Heating, cooling, and fan energy consumption shown below in Figure 37, Figure 38, and Figure 39, provides insight into how different control strategies balance different types of energy. Heating energy consumption is similar between the two strategies, but there are differences in cooling and fan energy consumption. Overall, CORE seems to be using less cooling energy, which aligns with the supply air temperature distribution discussed earlier. To maintain comfort in the building, CORE then has to use more airflow, with the same thermal load, which can be seen in Figure 39, especially around 70°F. However, the overall savings are achieved due to the difference in the magnitude of the two energy uses, where fan energy consumption is about half of the cooling energy consumption, around 70°F.

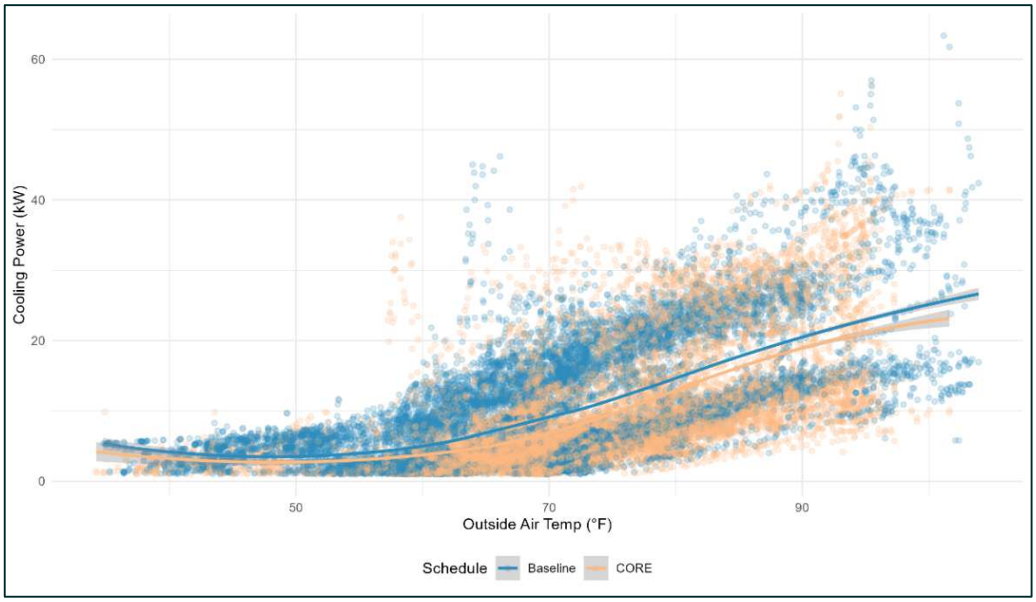

Figure 37: HVAC cooling energy consumption vs. OAT for combined RTU 2/3—MOB-2.

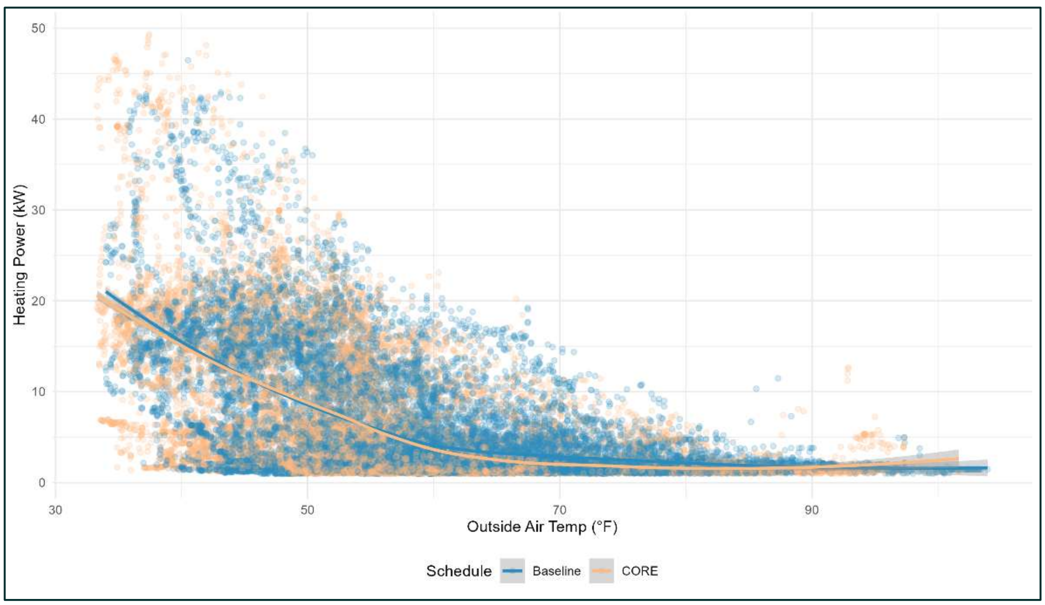

Figure 38: HVAC heating energy consumption vs. OAT for combined RTU 2/3—MOB-2.

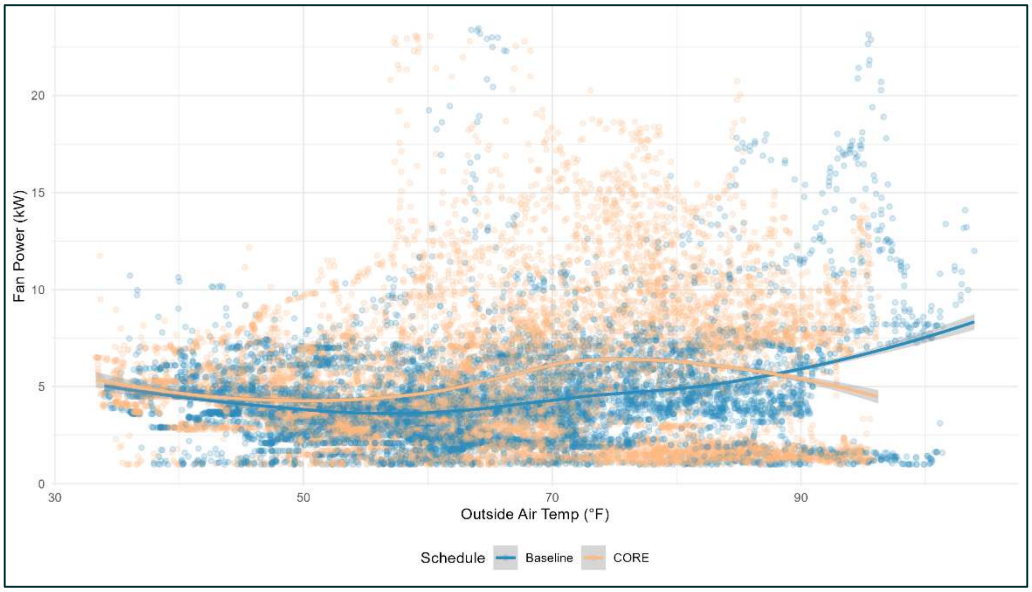

Figure 39: HVAC fan energy consumption vs. OAT for combined RTU 1 and RTU 2/3—MOB-2.

Similarly, Figure 40 shows the total heating, cooling, and fan cost for Baseline and CORE for the combined RTU 1 and RTU 2/3. Through the entire OAT range, CORE HVAC cost is lower than Baseline since CORE is always moving to setpoints that have the minimum HVAC cost.

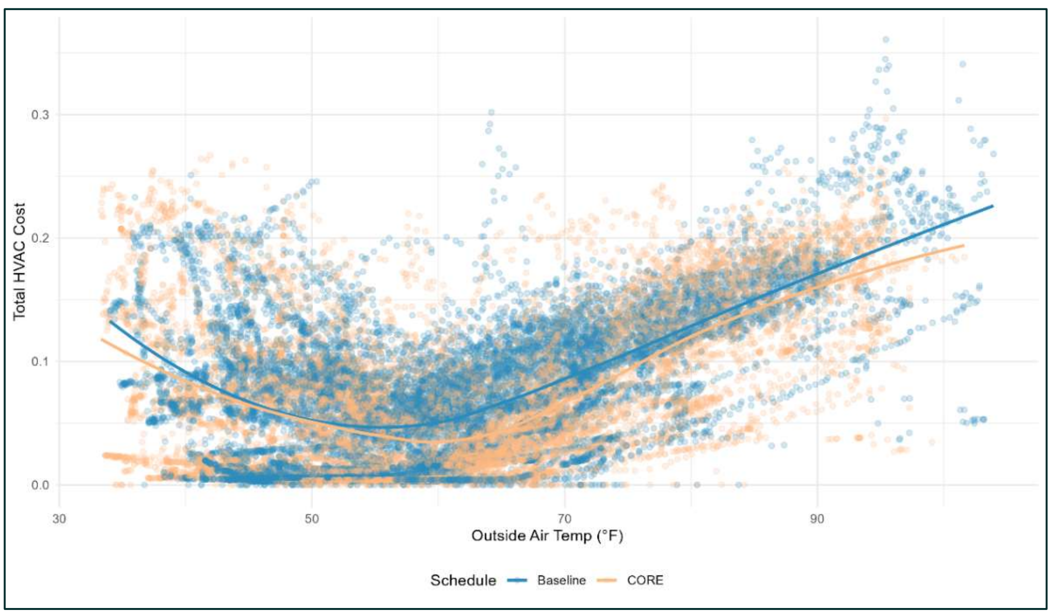

Figure 40: HVAC heating, cooling, and fan cost vs. OAT for combined RTU 1 and RTU 2/3—MOB-2.

Figure 41, Figure 42, and Figure 43, show the cooling, heating, and fan cost comparison between Baseline and CORE for combined RTU 1 and RTU 2/3. Like energy, cost is also balanced between cooling and fan, whereas the heating cost is similar between the two strategies.

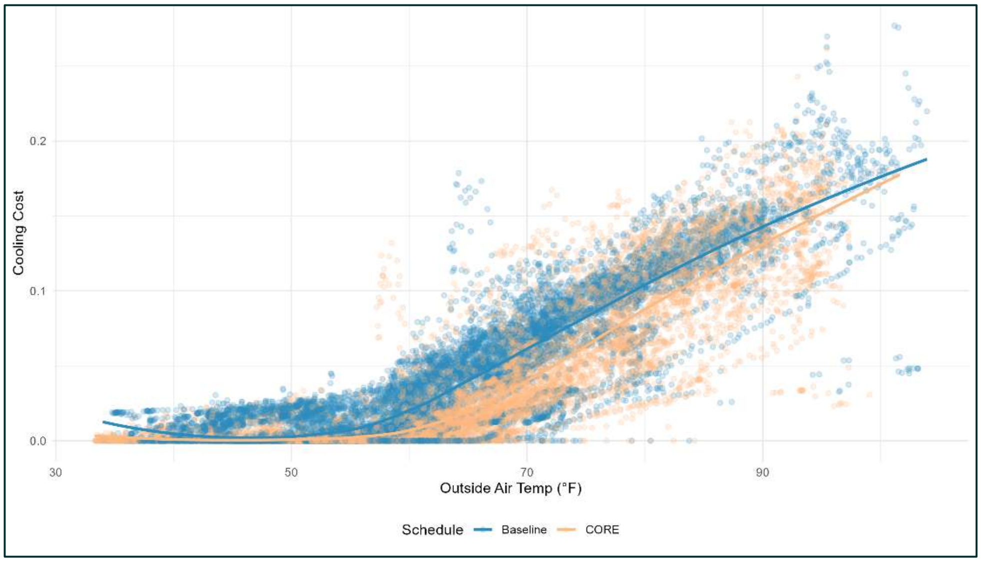

Figure 41: HVAC cooling cost vs. OAT for combined RTU 1 and RTU 2/3—MOB-2.

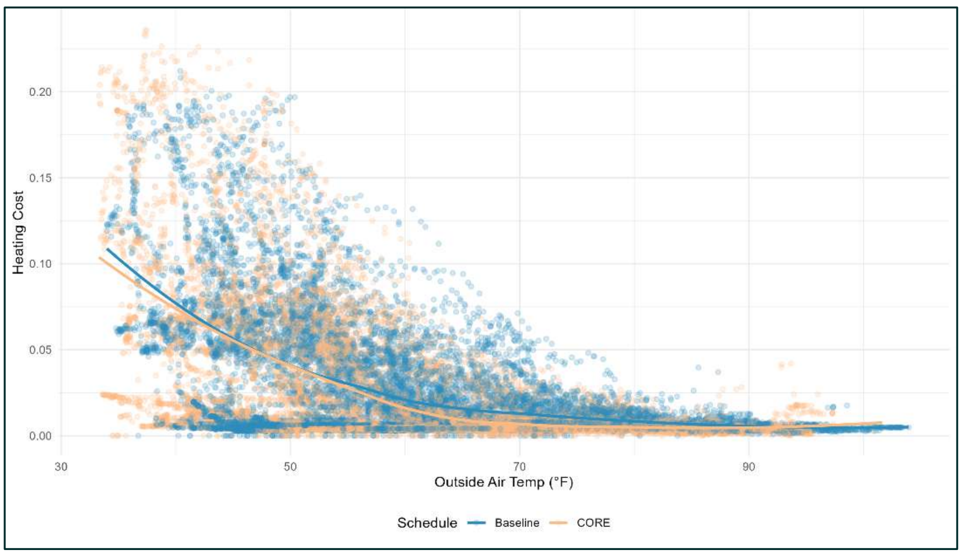

Figure 42: HVAC heating cost vs. OAT for combined RTU 1 and RTU 2/3—MOB-2.

Figure 43: HVAC fan cost vs. OAT for combined RTU 1 and RTU 2/3—MOB-2.

Figure 44 shows the total heating, cooling, and fan power density against time-of-day, split monthly. General variation of energy use during the day is similar across months, where midday OAT is at its highest and requires more cooling power. Similarly, power consumption increases for warmer months like June vs. mild weather months like March or April due to increased cooling demand.

The month of February has limited data of just seven days, and the distribution looks somewhat different to the other months. The team expects colder months to have higher energy consumption but more data is required for colder months—December to February—to determine if there is a significant difference in the energy use pattern.

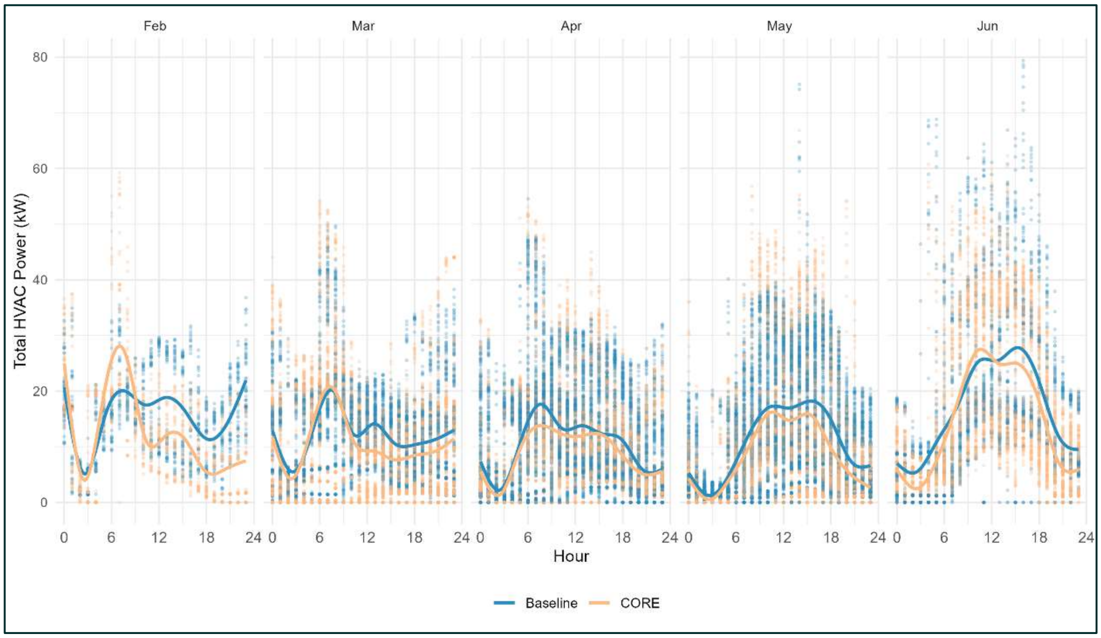

Figure 44: Total heating, cooling, and fan power distribution with time-of-day for different months—MOB-2.

Figure 45 below illustrates the heating, cooling, and fan energy cost distribution with time-of-day for each month. Similar to the energy distribution, warmer months, e.g., June, incur higher energy cost due to increased cooling demand.

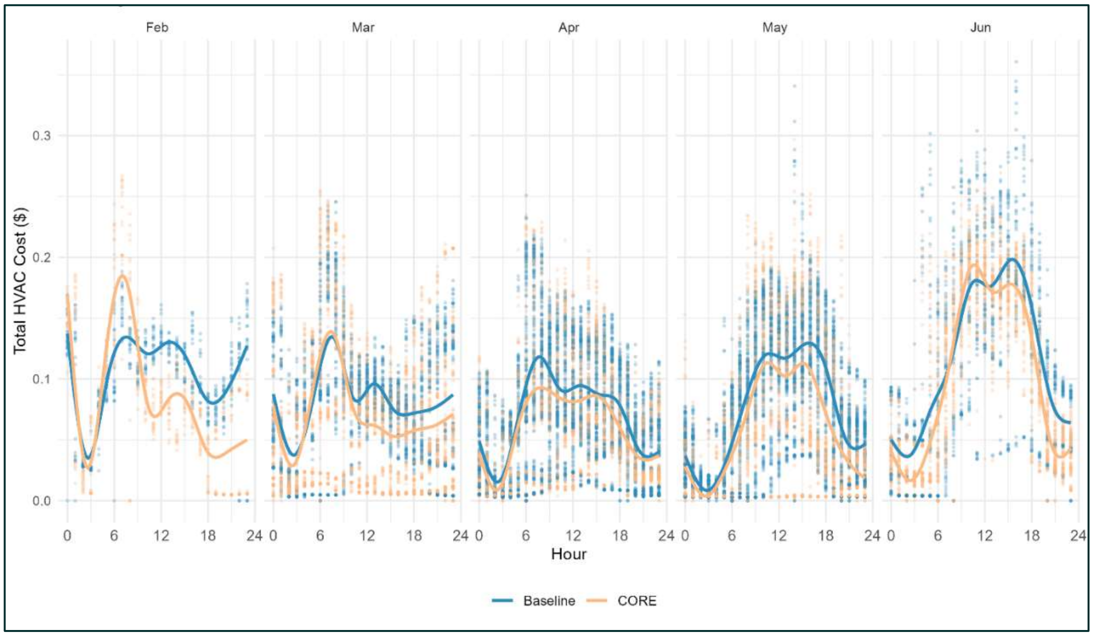

Figure 45: Total heating, cooling, and fan cost distribution with time-of-day for different months—MOB-2.

ACTUAL ENERGY AND COST DIFFERENCE

Figure 46 shows the daily energy total heating, cooling, and fan energy consumption distribution between Baseline and CORE. Daily mean energy consumption is similar between strategies with CORE consuming only two percent less energy on average. The p-value is much higher than 0.1, indicating that there is insufficient evidence to claim that CORE uses less energy than Baseline.

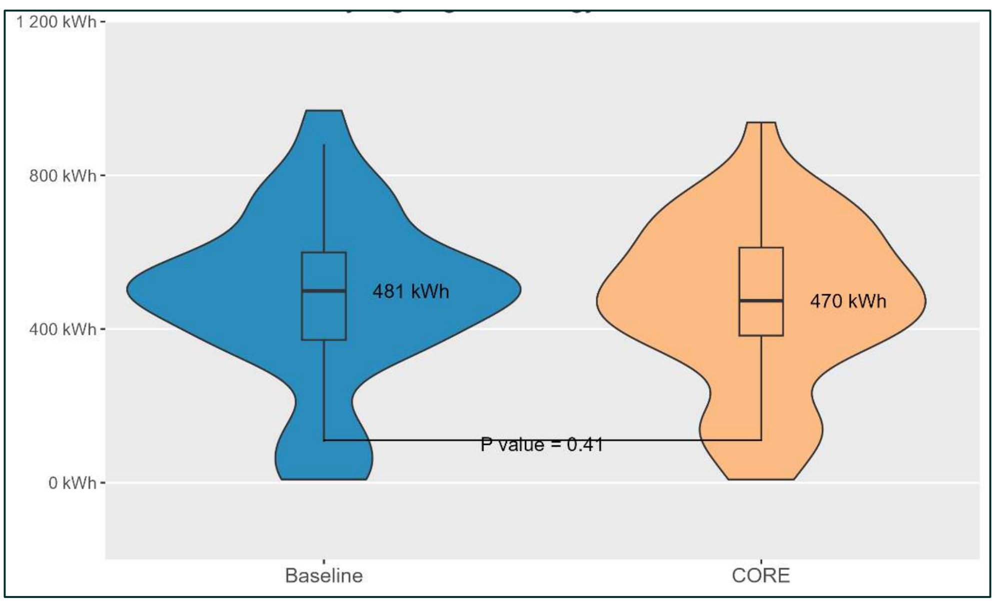

Figure 46: Daily total heating, cooling, and fan energy consumption distribution comparison between

Baseline and CORE—MOB-2.

Similarly, Figure 47 shows the energy cost distribution comparison between Baseline and CORE. CORE has 10 percent less cost with a better p-value. This comparison includes data for 40 days each, and the team believes the statistical significance will improve as more data is gathered.

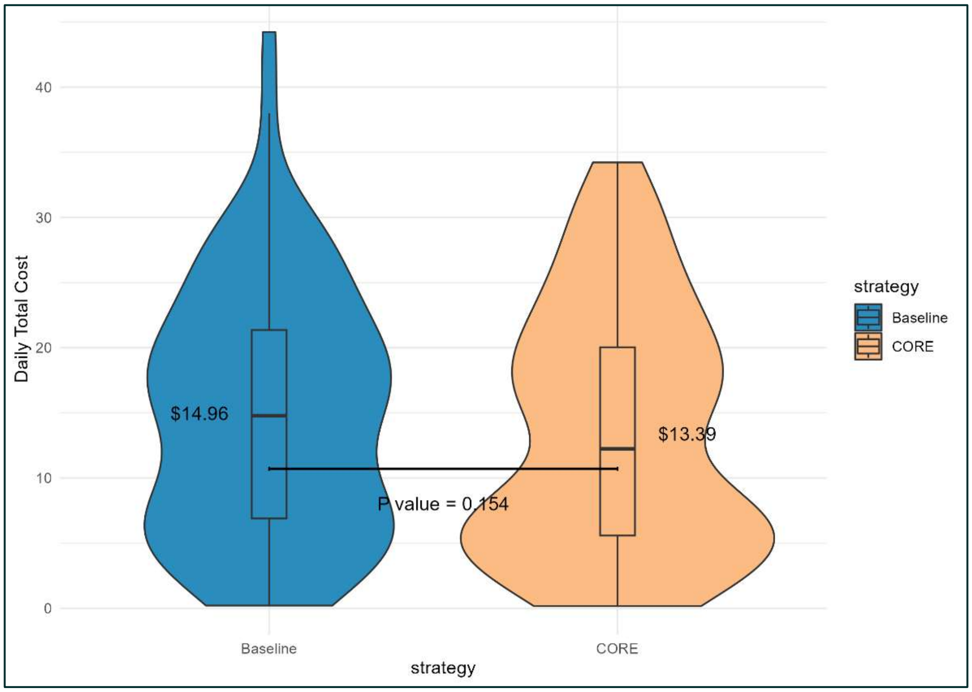

Figure 47: Daily total heating, cooling, and fan energy cost distribution comparison between Baseline and CORE—MOB-2.

NORMALIZED ENERGY USE IMPACT

The team used the TOWT method to develop regression models for Baseline and CORE energy consumption data. For MOB-2, total HVAC metered energy consumption was not available to the team, so the figures below focus on heating, cooling, and fan energy consumption.

Figure 48 shows the distribution of actual measured energy and the model-predicted energy for the same outdoor air temperature and time of week conditions against the outdoor air temperature as well as the difference between the actual and model-predicted (residual) for Baseline data. Similarly, Figure 49 shows the distribution for CORE data. The models Baseline and CORE seem to follow the measured energy trend fairly well, leading to a relatively high R^2 of 0.78 and 0.74, as can be seen in Table 11, low CVRMSE percent and NMBE percent values. This suggests that the TOWT model can predict the total heating, cooling, and fan energy consumption well for both Baseline and CORE. Furthermore, the distribution follows the expected trend where energy consumption is higher at high and low OATs and low at moderate OATs.

Figure 48: Measured energy vs. model-predicted energy for Baseline conditions vs. OAT (top) and residuals (actual minus modeled) vs. OAT (bottom)—MOB-2.

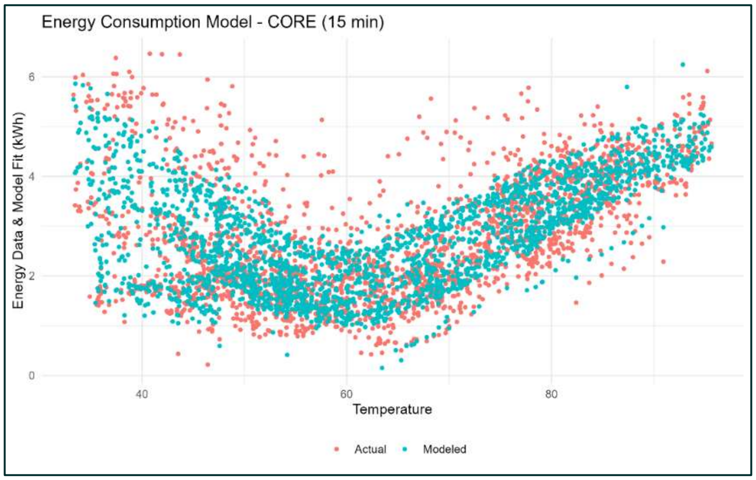

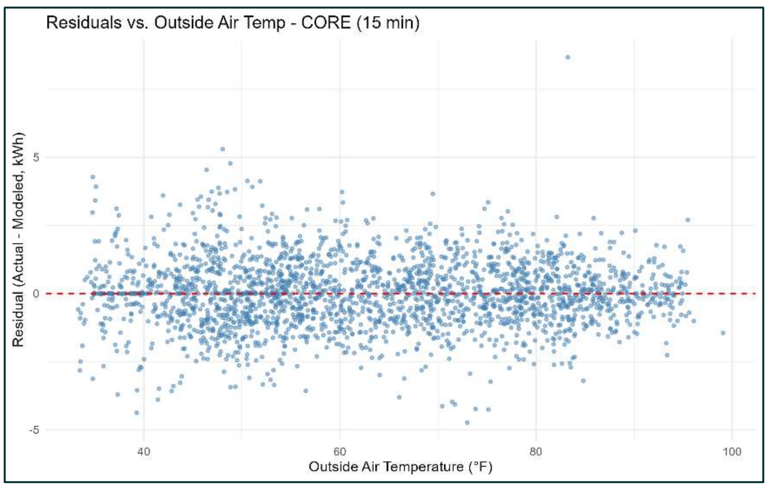

Figure 49: Measured energy vs. model-predicted energy for CORE conditions vs. OAT (top) and residuals (actual minus modeled) vs. OAT (bottom)——MOB-2.

The team used these TOWT models and normalized annual weather data from TMY3 to calculate the predicted normalized energy for a year. This data is summarized in Figure 50 below, where monthly energy is shown for both Baseline and CORE. Summer months have the highest energy consumption, followed by winter months, where shoulder seasons have less energy consumption than expected.

For January and December, the model predicts that CORE would consume more energy. However, the data used to train the model did not include data from these months, and so the prediction for these months should be re-evaluated by adding more data to the training dataset. For other months, CORE is using less energy than Baseline, indicating CORE's ability to save energy across different weather conditions throughout the year.

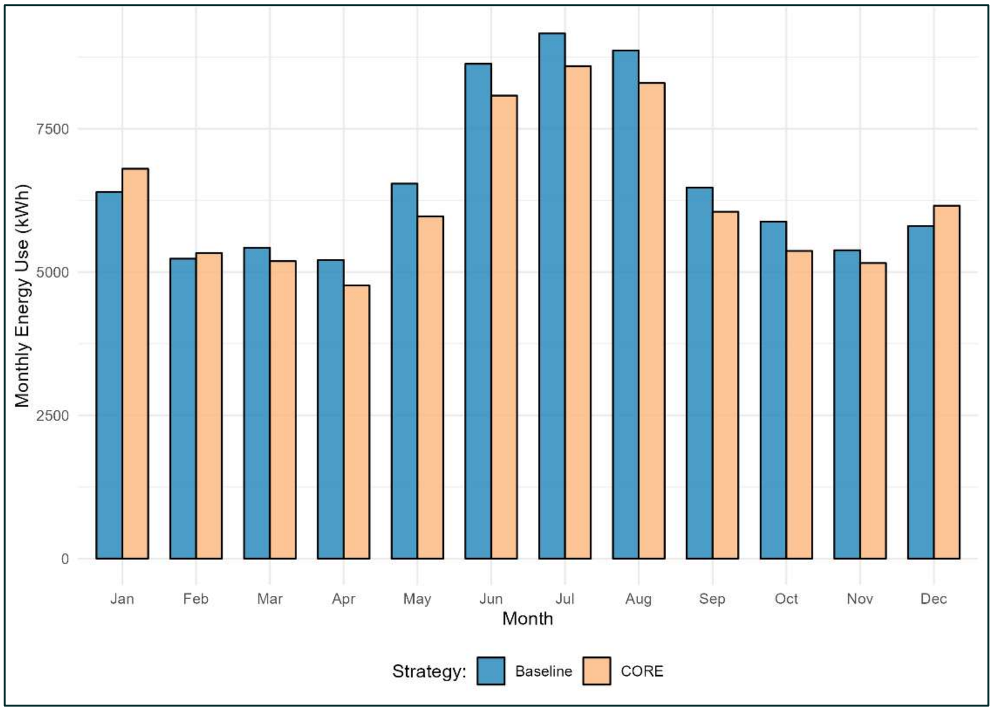

Figure 50: Predicted monthly heating, cooling, and fan energy consumption for Baseline and CORE for TMY3 weather data—MOB-2.

MOB-3

SAT CHARACTERIZATION

Figure 51 and Figure 52 below show the histogram of SAT setpoints for Baseline and CORE for the two AHUs at MOB-3. The temperature reset ranges for both Baseline and CORE are similar (55 to 65°F) with a morning temperature of 75°F.

Baseline strategy heavily favors low SAT setpoints, especially for AHU 1 where CORE operates at higher SAT setpoints. These lead to lower cooling and overall energy and cost for CORE.

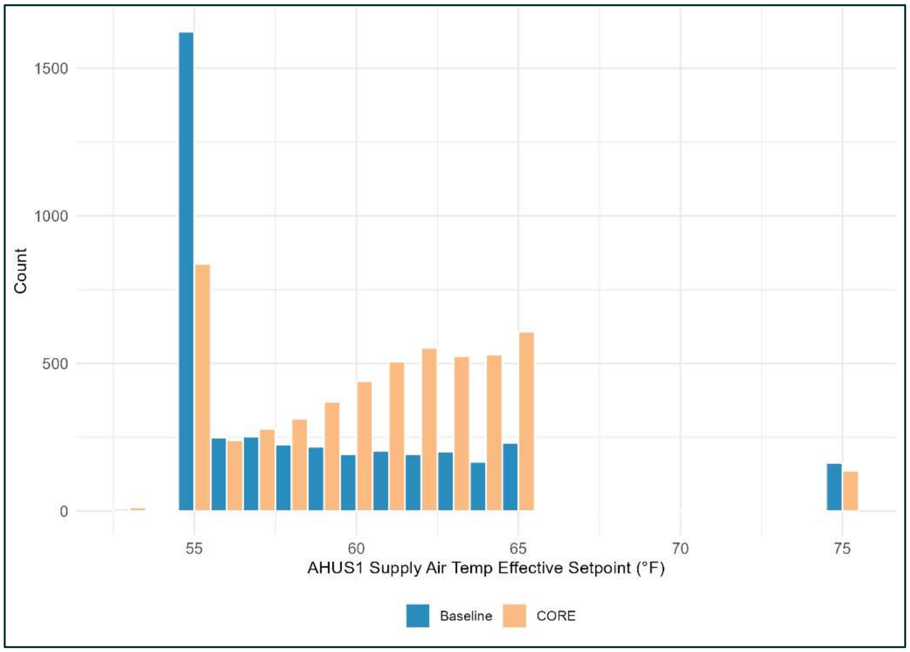

Figure 51: Supply air temperature setpoint distribution for Baseline CORE—AHU 1—MOB-3.

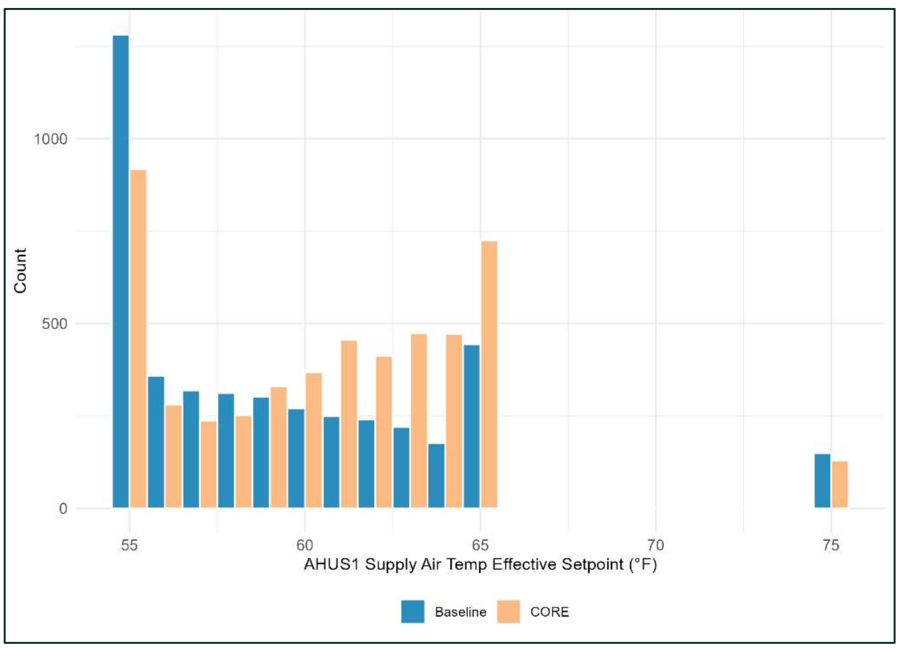

Figure 52: Supply air temperature setpoint distribution for Baseline and CORE—AHU 2—MOB-3.

Figure 53 and Figure 54 show the time-of-day distribution of SAT setpoint for Baseline and CORE for the two AHUs. In both cases, CORE uses higher SAT setpoints during most of the day. Towards the end of the when the cooling load is low, CORE moves to slightly lower SAT setpoints while likely reducing the air flow rates to reduce fan energy.

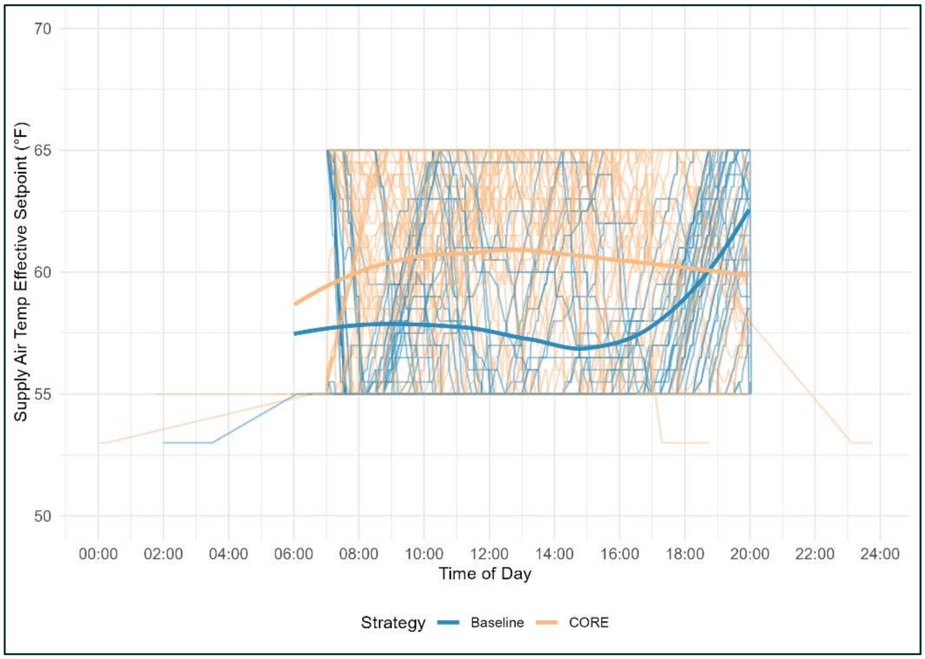

Figure 53: Supply air temperature distribution by time of day for AHU 1—MOB-3.

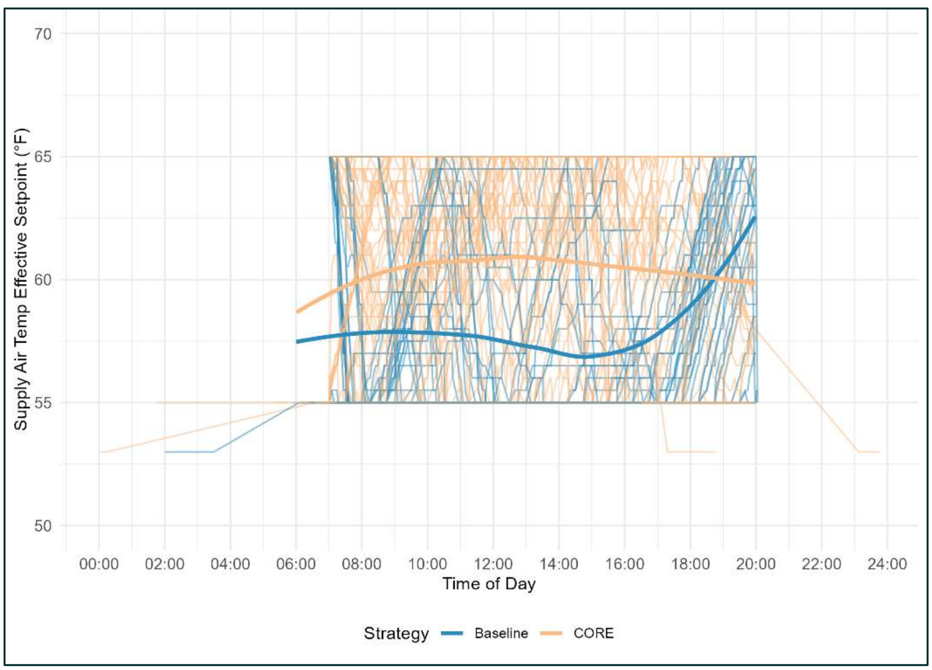

Figure 54: Supply air temperature distribution by time of day for AHU 2—MOB-3.

ENERGY AND COST CHARACTERIZATION

Figure 55 below shows the heating, cooling, and fan power consumption for combined AHU 1 and AHU 2 against the OAT. As expected, there is higher energy use at higher OATs due to higher cooling energy requirements.

Energy consumption trends for Baseline and CORE are similarly distributed, with CORE spending less energy throughout the range.

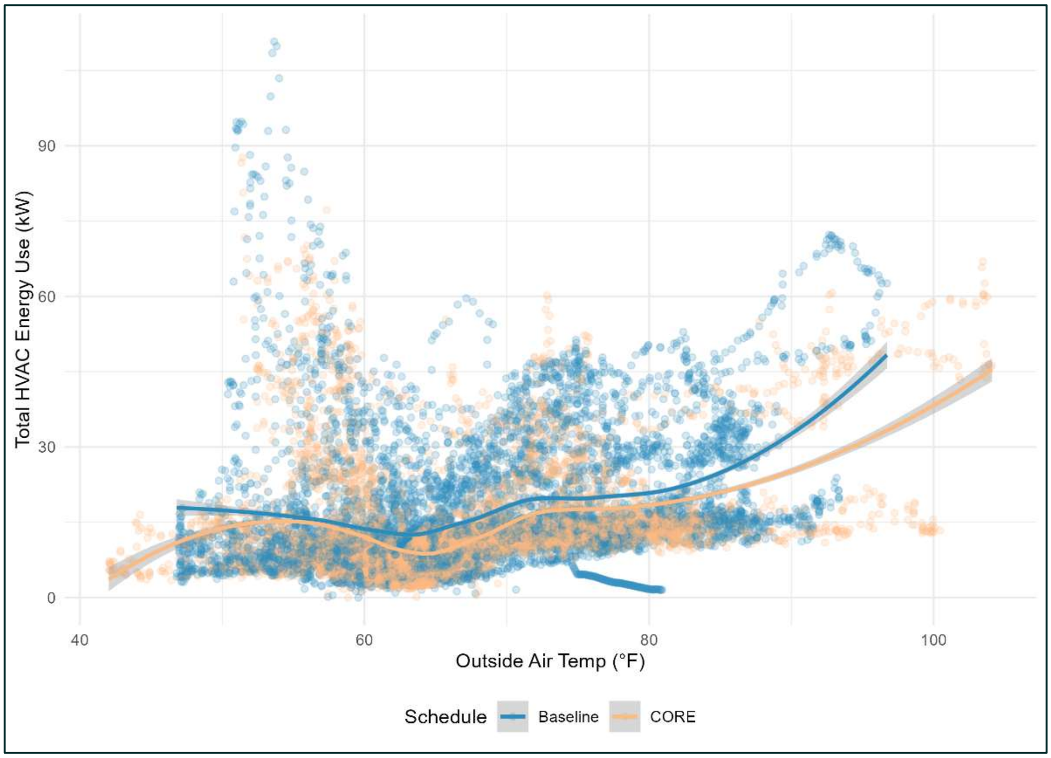

Figure 55: HVAC heating, cooling, and fan energy consumption vs. OAT for combined AHU 1 and AHU 2—MOB-3.

Looking at cooling, heating, and fan energy consumption with OAT illustrated in Figure 56, Figure 57, and Figure 58, cooling power seems to be lower and fan power seems to be higher for CORE, while heating power is very similar compared to Baseline. This is as expected since higher temperature as shown earlier, corresponds to lower cooling energy but higher fan energy due to the need to meet the comfort requirement within the same thermal load.

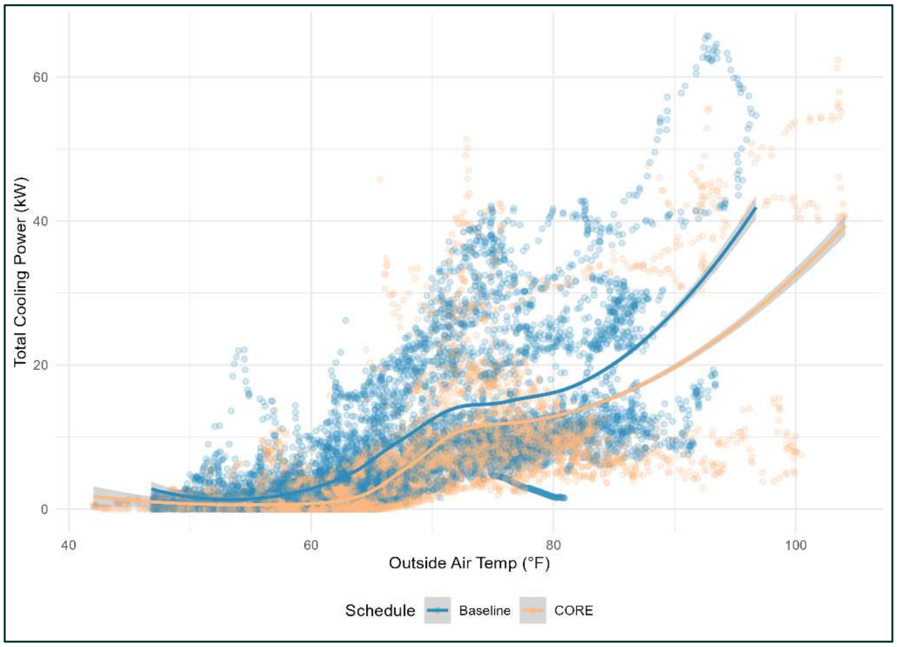

Figure 56: HVAC cooling energy consumption vs. OAT for combined AHU1 and AHU2—MOB-3.

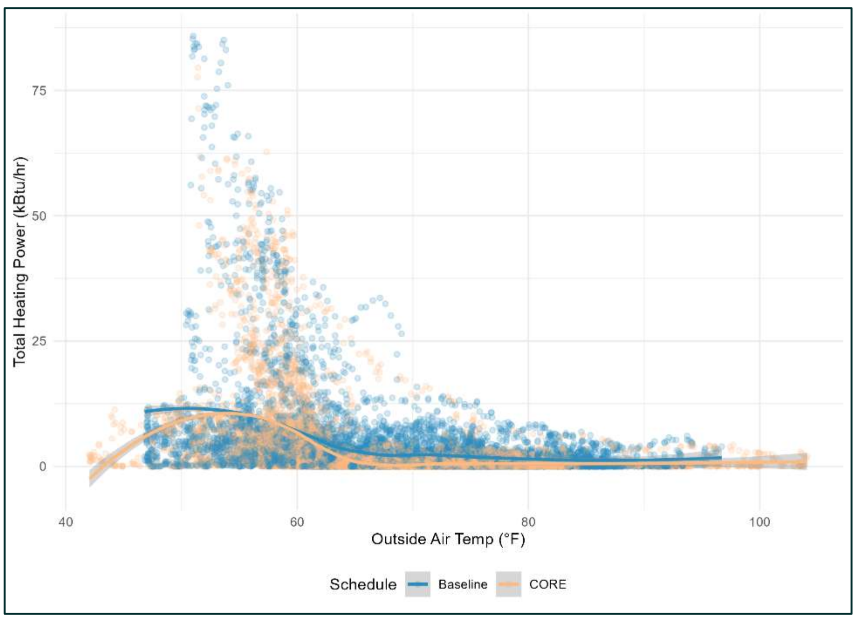

Figure 57: HVAC heating energy consumption vs. OAT for combined AHU1 and AHU2—MOB-3.

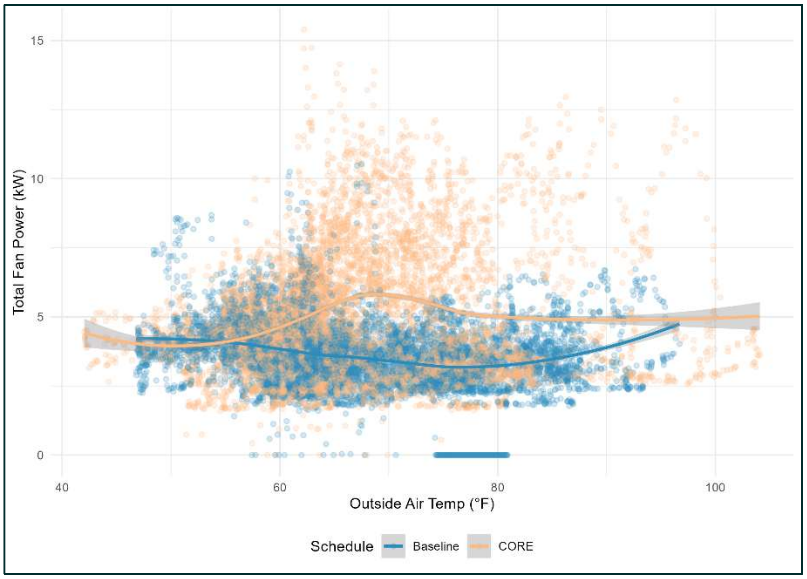

Figure 58: HVAC fan energy consumption vs. OAT for combined AHU1 and AHU2—MOB-3.

Similarly, Figure 59 shows the total heating, cooling, and fan cost for Baseline and CORE for the combined AHU 1 and AHU 2. Through the entire OAT range, CORE HVAC cost is lower than the Baseline since CORE is always moving to setpoints that have the minimum HVAC cost.

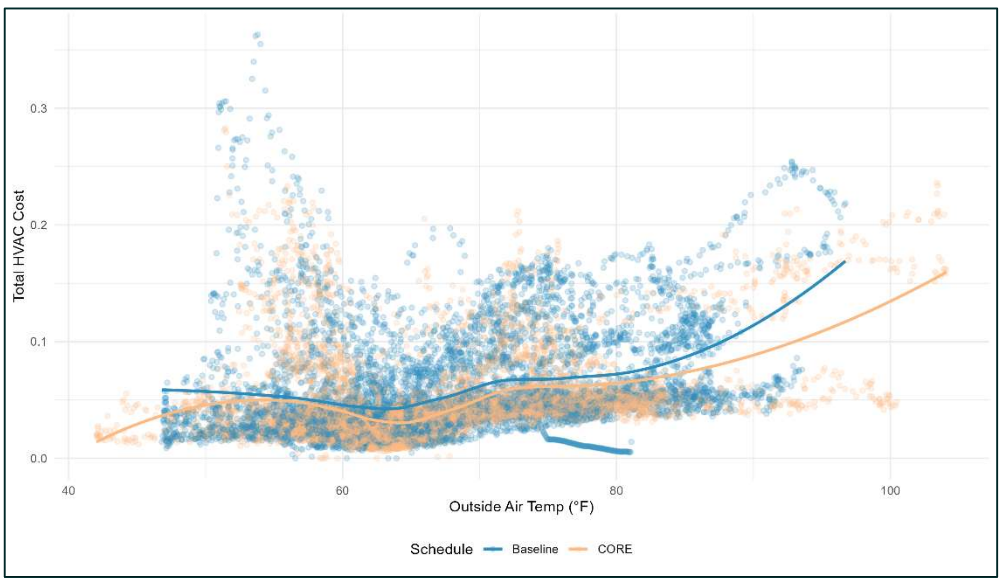

Figure 59: Total heating, cooling, and fan energy cost for combined AHU 1 and AHU 2—MOB-3.

Figure 60, Figure 61, and Figure 62 show the cooling, heating, and fan cost comparison between Baseline and CORE for combined AHU 1 and AHU 2. Similar to energy, cost is also balanced between cooling and fan where the heating cost is only slightly lower in CORE compared to Baseline.

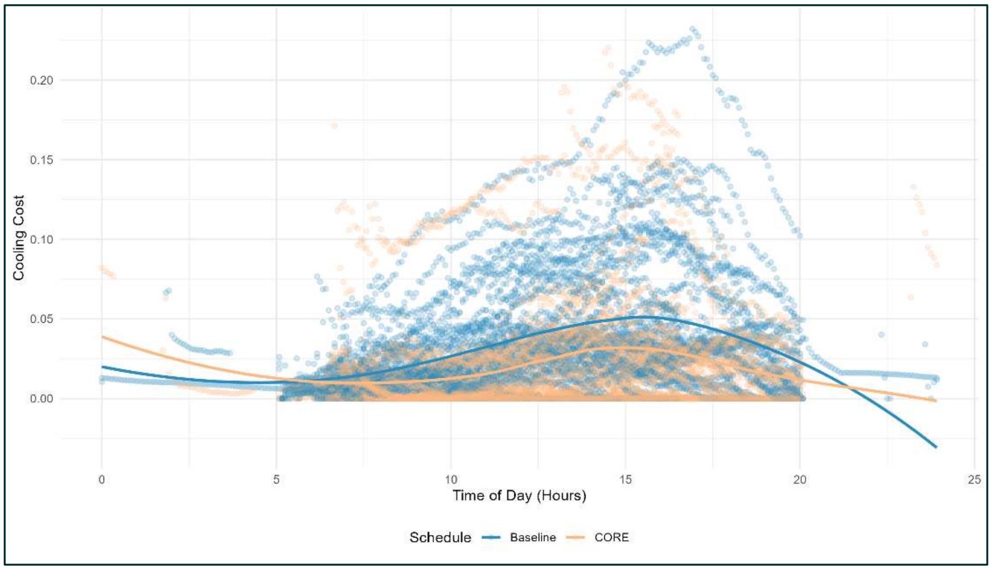

Figure 60: HVAC cooling cost vs. OAT for combined AHU 1 and AHU 2—MOB-3.

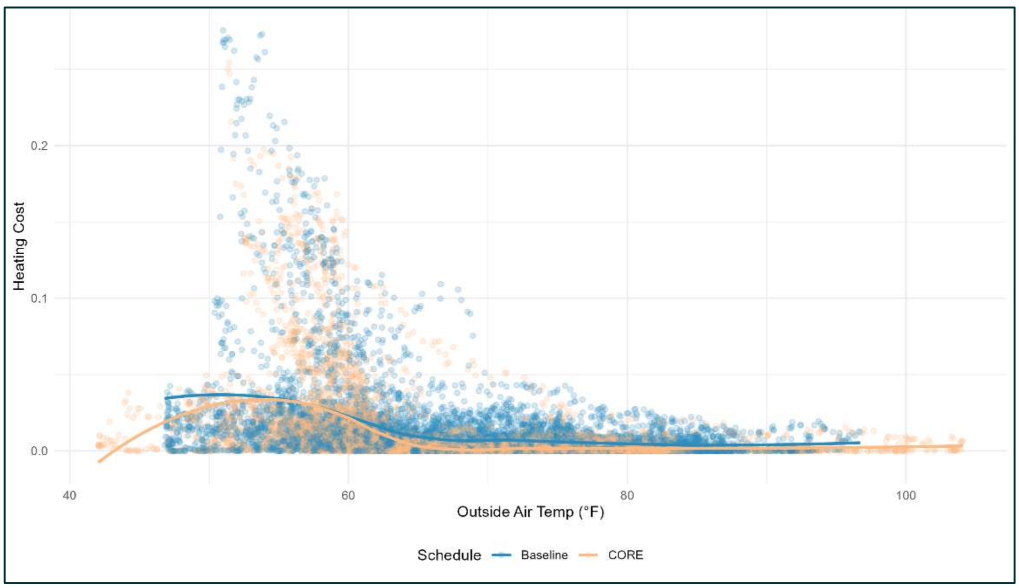

Figure 61: HVAC heating cost vs. OAT for combined AHU 1 and AHU 2—MOB-3.

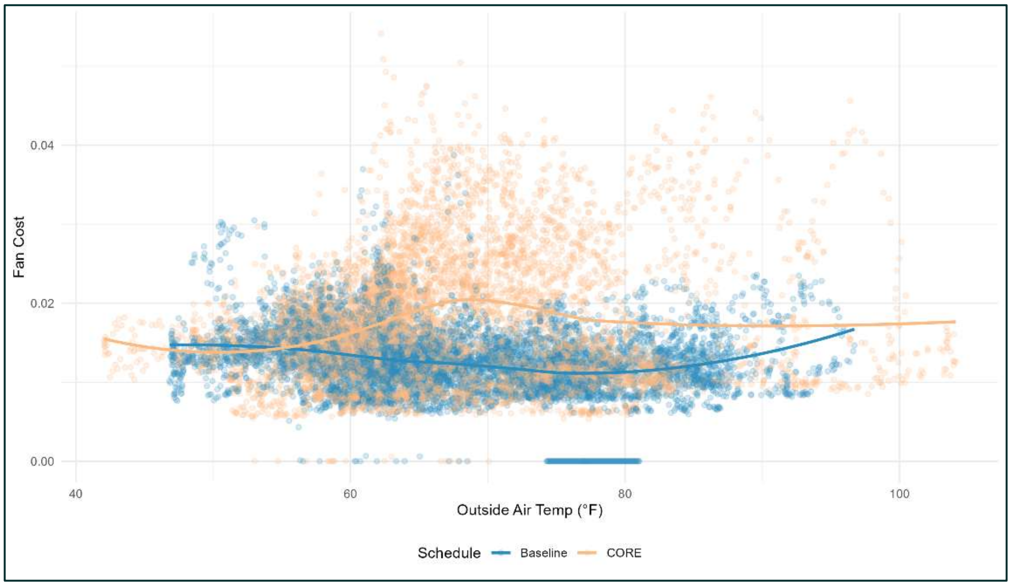

Figure 62: HVAC fan cost vs. OAT for combined AHU 1 and AHU 2—MOB-3.

Figure 63 below shows the total heating, cooling, and fan power density against time of day, split monthly. General variation of energy use during the day is similar across months, where midday OAT is at its highest, needing more cooling power. But unlike MOB-2, power consumption remains similar for June compared to mild weather months such as April and May.

Limited data from February (only seven days) are available and the energy usage seems to be higher compared to other months. This could be due to February being cold but more data from colder months is required to generalize this conclusion.

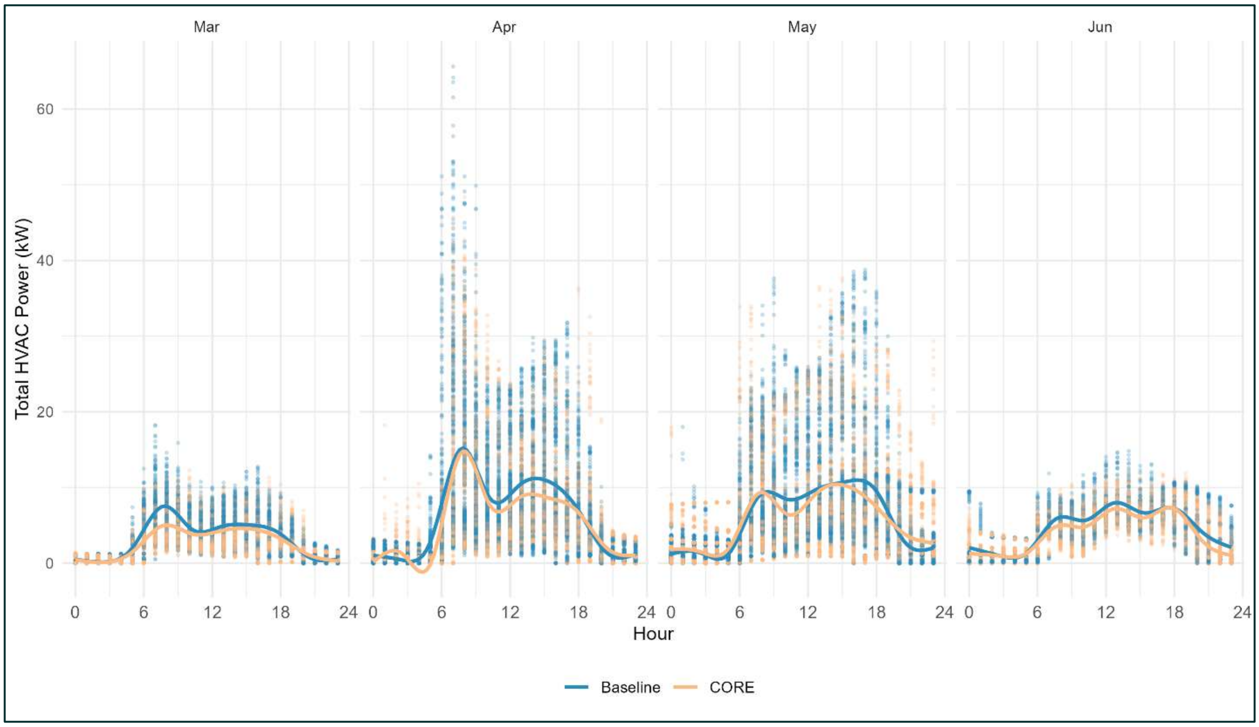

Figure 63: Total heating, cooling, and fan power distribution by time of day for different months—MOB-3.

Figure 64 shows similar density plots for total heating cooling and fan energy cost which follows similarly to power distributions. February, which has limited data, shows the largest difference between CORE and Baseline. However, CORE is consistently lower than the Baseline throughout the day.

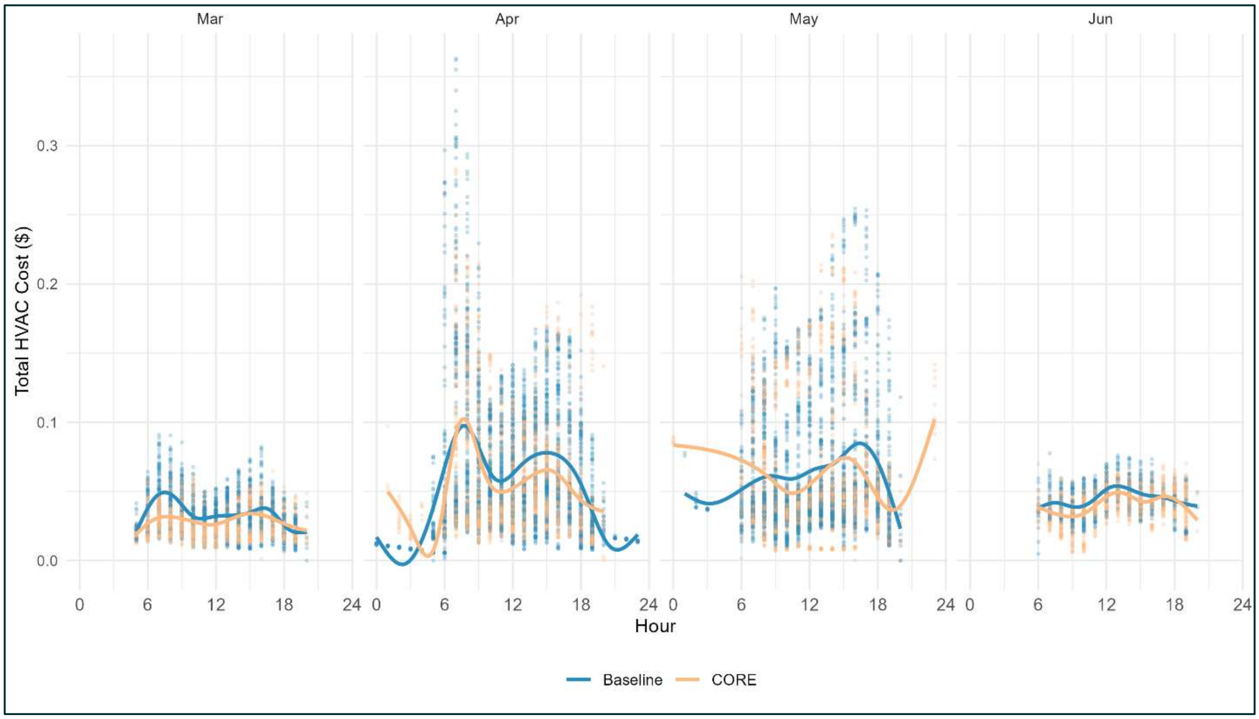

Figure 64: Total heating, cooling, and fan cost distribution by time of day for different months—MOB-3.

ACTUAL ENERGY AND COST DIFFERENCE

MOB-3 has submeters that allowed the research team to obtain HVAC electricity and natural gas usage. This data is used for analysis, shown below.

Figure 65 shows the daily combined electricity and gas energy consumption distribution between Baseline and CORE. Daily mean energy consumption is 15 percent less in CORE compared to the Baseline, with a p-value less than 0.1. This suggests there is evidence that CORE uses less energy.

Overall, this result shows CORE can save a significant amount of energy at the source level.

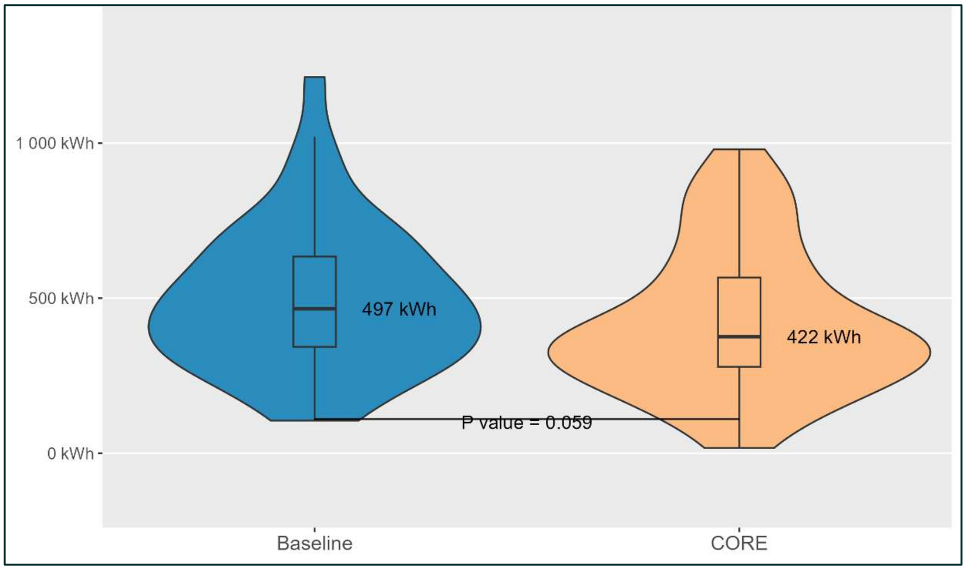

Figure 65: Daily HVAC electricity and gas consumption distribution comparison between Baseline and CORE—MOB-3.

Figure 66 shows the daily combined heating, cooling, and fan energy cost distribution between Baseline and CORE. CORE has 20 percent less cost with a p-value of less than 0.1, suggesting evidence for the claim that CORE saves energy cost.

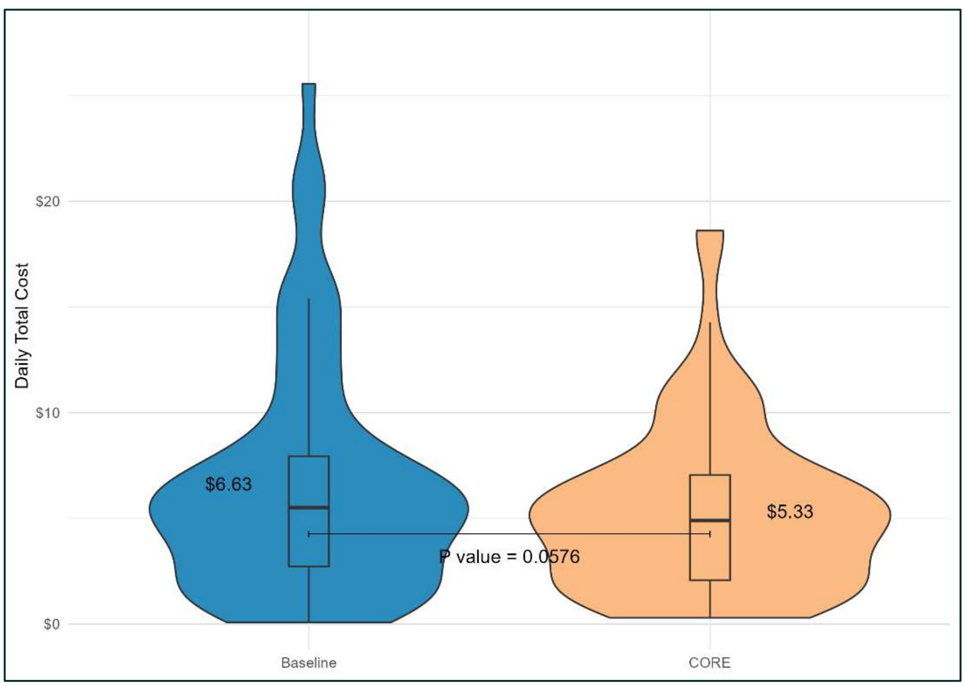

Figure 66: Daily HVAC Electricity and Gas energy cost distribution comparison between Baseline and CORE—MOB-3.

NORMALIZED ENERGY USE IMPACT

The team used the TOWT method to develop regression models for energy consumption for Baseline and CORE data. The team used the total electric and gas energy consumption data to develop the figures below.

Figure 67 below shows the distribution of actual measured energy and the model-predicted energy for the same outdoor air temperature and time of week conditions against the outdoor air temperature for Baseline data. Similarly, Figure 68 shows the distribution of CORE data. The models for Baseline and CORE seem to be following the measured energy trend fairly well, leading to relatively high R^2 of 0.72 and 0.68 as shown in Table 11, and relatively low CVRMSE percent and NMBE percent values. Even though these values do not meet the regression model thresholds, this suggests that the TOWT model is capable of reasonably predicting the total heating, cooling, and fan energy consumption for both Baseline and CORE. Furthermore, the distribution follows the expected trend where energy consumption is higher at high and low OATs and lower at moderate OATs.

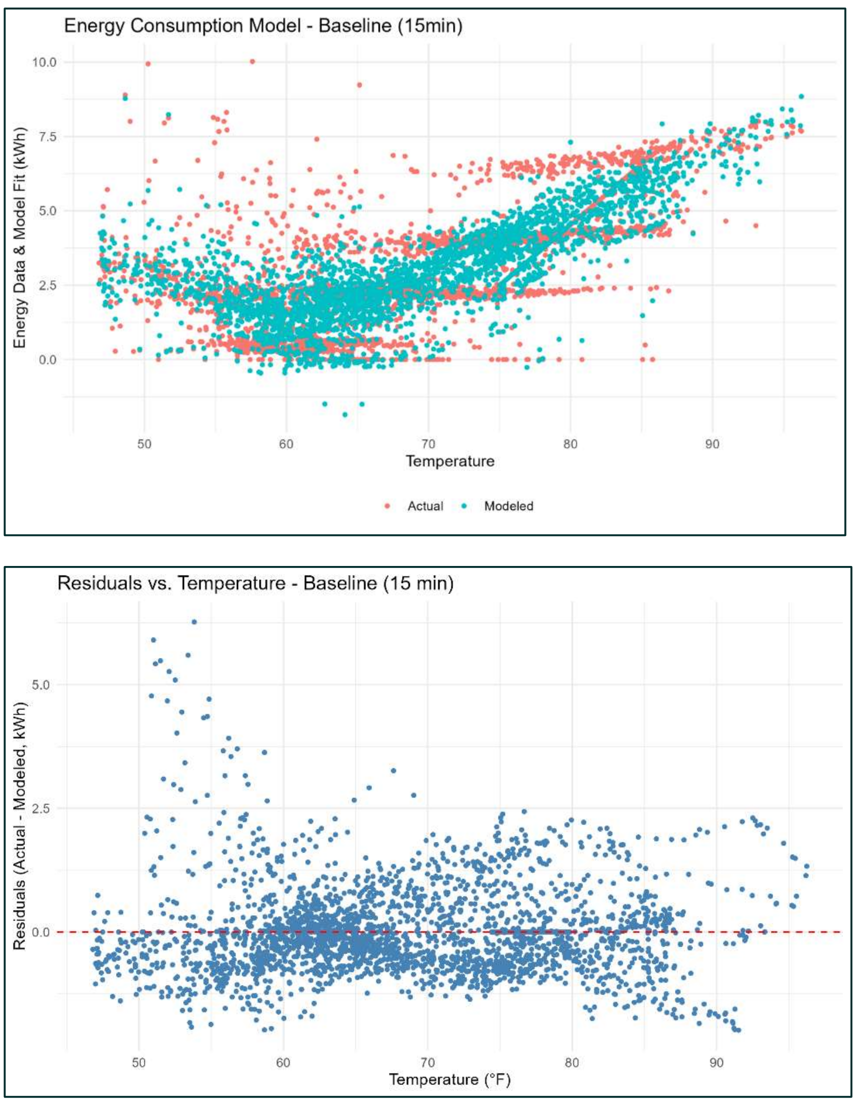

Figure 67: Measured energy vs. model-predicted energy for Baseline conditions vs. OAT (top) and Residuals

(actual minus modeled) vs. OAT (bottom)—MOB-3.

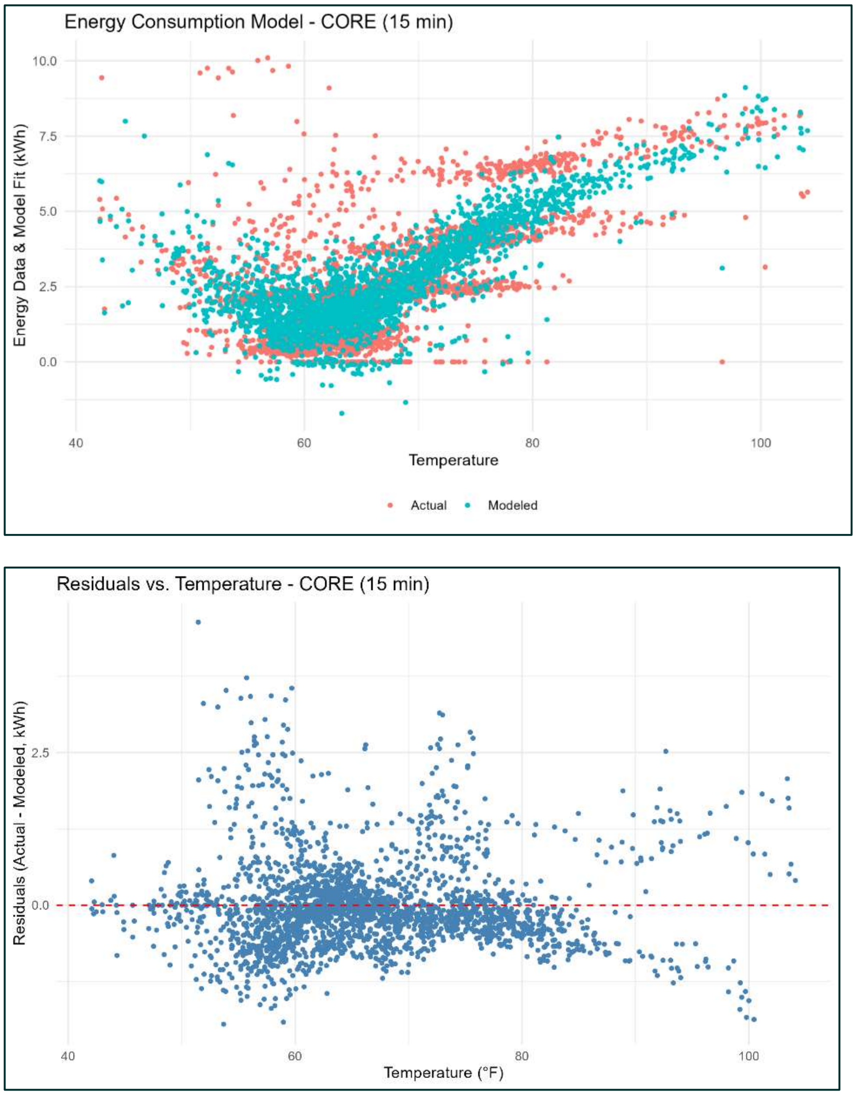

Figure 68: Measured energy vs. model-predicted energy for CORE conditions vs. OAT (top) and residuals (actual minus modeled) vs. OAT (bottom)—MOB-3.

The team used these TOWT models and normalized annual weather data from TMY3 to calculate the predicted normalized energy for a whole year. This data is summarized in Figure 69 below, where monthly combined electricity and gas HVAC energy is shown for both Baseline and CORE. Summer months have the highest energy consumption, and shoulder seasons have less energy consumption, as expected.

Predicted energy consumption between Baseline and CORE seems to be very close with some months even having higher energy consumption for CORE. However, it should be noted that the model training data did not include the coldest or warmest weather conditions.

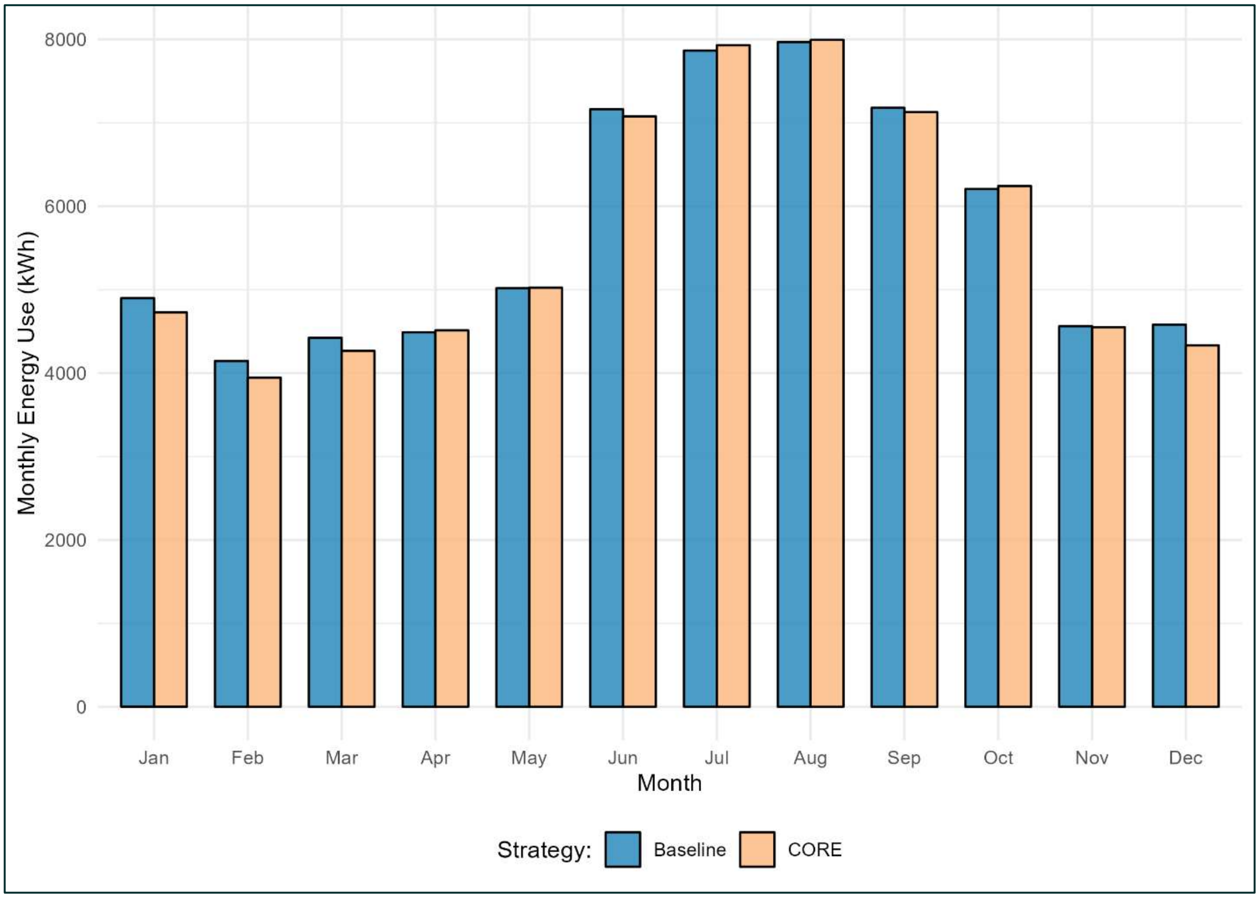

Figure 69: Predicted monthly HVAC electricity and gas energy consumption for Baseline and CORE—MOB-3.

Market Scalability Assessment

The team evaluated the CORE algorithm for its market scalability potential. The technical requirements for CORE are fairly minimal, and CORE applicability is based on HVAC system type and control capabilities. CORE requires a multi-zone VAV reheat HVAC system, a common system type for medium and large commercial buildings. In new construction, CORE could be included in the initial HVAC control system design, implementation, and commissioning as a highly efficient controls scheme.

A well-operating HVAC and control system is important for any improved control sequences to perform effectively. Deferred maintenance of mechanical sensors and equipment is a common barrier for advanced control implementation and should be considered as a required make-ready step and encouraged by facility managers and utility programs. CORE would be more cost-effective as part of a larger controls update or full retro-commissioning rather than as part of a standalone implementation. It will also be more cost-effective deploying in a large portfolio similar to how Altura implemented CORE at MOB-4. Costs to investigate issues, fix broken equipment, and control issues, as well as administrative costs such as developing contractor agreements, make CORE as a standalone retrofit not cost-effective. It may be more effective as an add-on to a routine vendor or contractor service. In the absence of a complete controls update or EBCx, at a minimum, the following should be true for a candidate building to ensure a well-operating system:

- A functioning air-side economizer
- VAV box minimums are at or near ventilation minimums
- SAT meets setpoint
- Zone space temperature deadbands are at least 4°F
- Rogue zones are actively managed.
- Modern control system with direct digital controls down to the zone terminal units.

One way to potentially get around some of the existing control system usability issues and decrease CORE implementation costs would be to have the CORE algorithm in an accessible programming library. Manufacturers could add CORE to their HVAC software and provide it to their control installers, and then a controls contractor would configure CORE for a specific site. The research team is providing CORE programming in a Python script that may be useful for some manufacturers, contractors, and vendors.

CORE could be marketed in a few ways for wider adoption:

- An energy conservation measure for EBCx processes, having the EBCx team evaluate the HVAC system for CORE implementation with potential savings tied to it.
- As a part of a larger control systems retrofit.

Lastly, a simplified CORE algorithm would be easier to implement and could get around existing systems' issues with data availability or communication bandwidth restrictions. This could limit the potential for misinterpretation of the algorithm and limit the number of potential issues from programming and point mapping.

Stakeholder Feedback

The team distributed the findings from the draft report with stakeholders including controls contractors, BAS manufacturers, consultants, and controls designers. We received feedback from Altura Associate (controls contractor) and a BAS manufacturer. Most of the comments were minor and highlighted the below.

- The BAS manufacturer highlighted the importance of data transfer pipeline for CORE operation. This can be especially challenging in older buildings even with supervisory controls layer since the bandwidth and speed of equipment data transfer can mean the supervisory layer does not receive the latest BAS data for the CORE calculation. The simplified CORE algorithm (CORE 2) that does not use any zone level data having good performance in energy simulation results was highlighted as a promising option to reduce the data traffic.
- The BAS manufacturer also highlighted the need to consider heating requests in buildings in colder climates where some zones in the building could be cooling-only.
- Both reviewers supported the benefits of CORE to optimize overall energy consumption and cost considering heating, cooling and fan energy use.

Recommendations

In this field demonstration project, the team refined the previously developed CORE control method to be adoptable for a wide range of buildings in California, conducted a detailed energy simulation analysis to evaluate energy savings expected from CORE compared to industry common practice and industry best practice approaches to control SAT, and implemented CORE in three commercial buildings to evaluate energy and energy cost savings.

The energy simulation study highlighted that CORE consistently performed better than all other strategies across different building load conditions and climate zones. This suggests that CORE can have a statewide impact on saving energy and reducing energy costs.

The field demonstration in the three medical office buildings highlighted the importance of existing building equipment and control system operation to CORE functionality and its ability to optimize energy cost. The project team recommends a thorough evaluation of the existing HVAC system and BAS controls, and a process to identify and resolve issues before implementing any new controls retrofits. The team also identified the importance of supporting the implementers on CORE programming and deployment in buildings so that the intention of CORE is correctly interpreted and programmed. For this purpose, the team believes that CORE may be better packaged within a programming library offered by BAS manufacturers so that it is better integrated with the other controls rather than implementers programming and deploying CORE as a standalone measure.

For CORE to be consistently implemented, there will need to be standards, processes, and case studies developed to include CORE in various types of projects. Further efforts are recommended to:

- Develop pre-programming in coordination with control manufacturers.
- Define control system functionality and accessibility requirements for implementation in new and existing buildings.
- Define minimum HVAC operational requirements to fully utilize CORE.
- Ease implementation path through controls retrofit and EBCx focused standards and guidelines.

Easing the implementation path for engineers, controls manufacturers and installers, and facility management/operations through these recommendations can help CORE be more effective and installed in more locations.

References

ASHRAE. 2024. *Guideline 36-2024 - High Performance Sequences of operation for HVAC systmes.* Guideline, ASHRAE. https://www.techstreet.com/ashrae/standards/guideline-36-2021-high-performance-sequences-of-operation-for-hvac-systems?product_id=2229690.

Bucaneg, Haile, and RJ Wichert. 2022. *2022 Nonresidential And Multifamily Alternative Calculation Method Reference Manual: For the 2022 Building Energy Efficiency Standards Title 24, Part 6, and Associated Administrative Regulations in Part 1.* California Energy Commission.

n.d. *CalNEXT.* Accessed 10 18, 2023. https://calnext.com/.

Granderson, J, P Gruendling, C Torok, P Jacobs, and N Gandhi. 2019. *Site-Level NMEC Technical Guidance: Program M&V Plans Utilizing Normalized Metered Energy Consumption Savings Estimation.* LBNL. https://www.cpuc.ca.gov/-/media/cpuc-website/files/legacyfiles/l/6442463695-lbnl-nmec-techguidance-01072020.pdf.

Granderson, Jessica, Guanjing Lin, Rupam Singla, Samuel Fernandes, and Samir Touzani. 2018. "Field evaluation of performance of HVAC optimization system in commercial buildings." *Energy and Buildings* 173: 577-586.

Mathieu, Johanna L., Phillip N. Price, Sila Kiliccote, and Mary Ann Piette. 2011. "Quantifying Changes in Building Electricity Use, with Application to Demand Response ." *IEEE Transactions on Smart Grid.* IEEE.

PG&E. 2024. *Electric Schedule B-10.* 05 01. Accessed 05 01, 2024. https://www.pge.com/tariffs/assets/pdf/tariffbook/ELEC_SCHEDS_B-10.pdf.

—. 2024. *Gas Schedule G-NR2.* 05 01. Accessed 05 01, 2024. www.pge.com/tariffs/assets/pdf/tariffbook/GAS_SCHEDS_G-NR2.pdf.

Raferty, Paul, Aoyu Zou, Thomas Parkinson, Geoff Hancock. 2024. "Reliably estimating the impact of an active control strategy in a building." *Journal of Building Engineering.*

Raftery, Paul, Shuyang Li, Baihong Jin, Min Ting, Gwelen Paliaga, and Hwakong Cheng. 2018. "Evaluation of a cost-responsive supply air temperature reset strategy in an office building." *Energy and Buildings* 158: 356-370.

Appendix A: CORE Sequence of Operation

Multiple Zone VAV Air Handlers

This sequence excerpt applies to multiple zone VAV air handlers and is written to integrate with ASHRAE Guideline 36 High Performance Sequences of Operation for HVAC Systems. See Guideline 36 for explanation of Trim & Respond logic and for related sequences.

Supply Air Temperature Control

Control loop is enabled when the supply air fan is proven on and disabled and output set to zero otherwise.

Supply Air Temperature Setpoint

During Occupied Mode and Setup Mode: Setpoint shall be reset from Min_SAT (the lowest cooling supply air temperature setpoint) up to Max_SAT using Trim & Respond logic with the following parameters:

Variable	Value
Device	Supply Fan
SP_0	SP_{max}
SP_{min}	Min_SAT
SP_{max}	Max_SAT
T_d	10 minutes
T	5 minutes
I	2
R	See below
SP_{trim}	+0.5 °F
SP_{res}	-0.5 °F
$SP_{res-max}$	-2.0 °F

Requests:

If $R_{cool} > I$, where R_{cool} = Zone Cooling SAT Requests, then $R = R_{cool}$

Otherwise, $R = R_{cost}$ where

If $C_{lower} < C_{current}$ and $C_{lower} <= C_{higher}$, $R_{cost} = I + 2$ (decrease supply air temperature by SP_{res})

If $C_{higher} < C_{current}$ and $C_{higher} < C_{lower}$, $R_{cost} = 0$ (increase supply air temperature by SP_{trim})

Else, $R_{cost} = I + 1$ (no change in supply air temperature)

See Cost-Based Optimization section below for cost calculations: C_{lower}, C_{higher}, and $C_{current}$

During Cool-Down Mode: Setpoint shall be Min_SAT.

During Warm-Up and Setback Modes: Setpoint shall be 95 °F.

Cost-Based Optimization

This cost-based optimization approach is intended to apply to single-duct VAV reheat systems with DX cooling, chilled water, and/or hot water sources. It requires discharge air temperature (DAT) sensors at reheat terminals as well as airflow measurement at every VAV terminal. An alternate formulation can remove the DAT requirement, though this incurs an accuracy penalty. This sequence is not intended to apply to systems heating coils at the air handler. Logic shall be provided in the event of out-of-range measurements or non-numeric values due to device failure or calibration issues or communication loss and to protect against divide-by-zero calculation errors. Logic shall also be provided to limit excessive network traffic such as by limiting the update of values based on a change of value threshold or a set time interval (e.g. 30 seconds).Energy use measurement and estimates for this system shall be evaluated at the current SATs and at each of two alternate SAT setpoints (current SAT + SP_{trim} and current SAT + SP_{res})AirflowsMeasured airflows (in cfm) at current SAT shall be determined as follows:

Zone supply airflow Vz = airflow measured at each cooling-only and reheat VAV box

System supply airflow Vs = the sum of Vz values from all associated VAV boxes

Estimated airflows (in cfm) at alternate SATs shall be determined as follows:

For each reheat VAV zone *currently in cooling mode*: Estimated zone supply airflow Vz_alt = (Tz – Td) / (Tz – Td_alt) * Vz, where

Tz = Zone air temperature

Td = Discharge air temperature at zone terminal

Td_alt = Discharge air temperature at zone terminal at alternate SAT and is calculated as Td + (Ts_alt – Ts)

Vz = Zone supply airflow

Ts = Supply air temperature setpoint at air handler

Ts_alt = Alternate supply air temperature setpoint at air handler

Vz_alt shall be constrained to be no less than zone minimum airflow setpoint Vmin and no greater than zone maximum cooling airflow setpoint Vcool-max. If Tz is more than 1 °F greater than the zone cooling setpoint, the normal calculation shall be bypassed and Vz_alt set equal to Vcool-max.

For each cooling-only VAV zone *currently in cooling mode* (if no discharge air temperature sensor available): Estimated zone supply airflow Vz_alt = (Tz – Ts – T_{hg}) / (Tz – Ts_alt – T_{hg}) * Vz, where

Vz_alt shall be constrained to be no less than zone minimum airflow setpoint Vmin and no greater than zone maximum cooling airflow setpoint Vcool-max

T_{hg} = Estimated average temperature rise between air handler leaving air temperature and VAV box discharge air temperature, e.g., 2 °F.

For each zone in heating or deadband mode: Estimated zone supply airflow Vz_alt = Vz. Zone airflow does not directly change due to SAT adjustments when zone is in heating or deadband modes.

Estimated system airflow Vs_alt = the sum of Vz_alt values from all associated VAV boxes

Cooling coil energy rate (or power, in Btu/h) shall be estimated at current and alternate SATs based on a sensible heat balance across the cooling coil when the valve is open, as follows. Note that this approach does not directly account for latent cooling:

At current SAT: $P_{chw} = \max[0, 1.08 * (Tm - Ts + \Delta Tc) * Vs]$

At alternate SATs: $P_{chw\_alt} = \max[0, 1.08 * (Tm - Ts\_alt + \Delta Tc) * Vs\_alt]$

Where

Tm = Mixed air temperature at air handler. Where a measured mixed air temperature sensor is unavailable, an estimated value can be calculated using outdoor and return air temperature, and outdoor and supply airflow rates.

ΔTc is a temperature correction to account for fan heat, sensor drift, and/or passing control valves and is an exponential average equal to [k * (Ts − Tm) + (1 − k) * (ΔTc from last time step)] calculated during periods when the chilled water control valve has been closed for a minimum of 5 minutes and when airflow is proven. The value of ΔTc is fixed at its last value prior to the valve opening and for 5 minutes after closing. The exponential smoothing coefficient k is user-adjustable, with a default value between 0.01 and 0.001. A small value such as 0.01 can be used to ensure the smoothing can reflect the most recent changes.

P_{chw} and P_{chw_alt} are also smoothed exponentially with a user-adjustable smoothing coefficient of 0.5 to reduce compressor/chiller water valve cycling when operating near the point where a slightly cooler SAT begins to require mechanical cooling.

Zone reheat coil energy rate (or power, in Btu/h) shall be estimated at current and alternate SATs based on a sensible heat balance across each zone reheat coil when reheat control valve is open, as follows. Note that this approach does not directly account for waterside distribution losses. Alternatively, for modulating electric resistance reheat, coil energy rate can be directly estimated from heating output percentage and the nominal coil rating.:

At current SAT: P_{rh} = max[0, 1.08 * (Td − Ts − ΔTh) * Vz]

At alternate SATs: P_{rh_alt} = max[0, 1.08 * (Td − Ts_alt − ΔTh) * Vz]

Where ΔTh is a temperature correction to account for fan heat, duct gain, sensor drift, and/or passing control valves and is an exponential average equal to [k * (Td − Ts) + (1 − k) * (ΔTh from last time step)] calculated during periods when the reheat valve has been closed for a minimum of 5 minutes and when airflow is proven. The value of ΔTh is fixed at its last value prior to the valve re-opening and for 5 minutes after closing, value between 0.01 and 0.001. A small value such as 0.001 can be used to ensure the smoothing takes effect over a period of days. Note that Td and Vz are unchanged regardless of SAT adjustment. Total reheat coil energy rate P_{hhw} for this system shall be equal to the sum of P_{rh} from each associated zone, and evaluated for each SAT. Fan power (in kW) shall be determined for current and each alternate SAT as follows

At current SAT: fan power P_{fan} = power measured by variable speed drive or dedicated meter

At alternate SATs: estimated fan power $P_{fan\_alt} = P_{fan} * (Vs\_alt / Vs)^{2.5}$ based on affinity laws

Thermal energy conversions. Thermal energy use for each of the current and alternate SATs shall be converted to utility energy use as follows:

Cooling power P_c (kW) = P_{chw} * E_c * (1 ton /12000 Btu/h), where E_c is the chiller plant efficiency in units of kW/ton. This value may be a constant based on expected plant performance (e.g., 0.7 kW/ton), calculated based on performance curves, or may be replaced with real-time kW/ton measurements if power from all plant components and total ton measurements are available. For DX coils, assume 1-1.3 kW/ton for typical units.

Heating power P_h (Btu/h) = P_{hhw} / E_h, where E_h is the dimensionless boiler plant efficiency. This value may be a constant based on expected plant performance (e.g., 0.8), calculated based on performance curves, or may be replaced with real-time efficiency measurements if boiler gas use and total Btu/h measurements are available. For electric heating systems, a value of 1 can be used for resistance heating, and a higher value for heat pumps reflecting the coefficient of performance of the system (system specific, typically ranging from 2-4 in heating).

Energy cost calculations. HVAC energy cost calculations shall be evaluated for each of the current and alternate SATs ($C_{current}$, C_{lower}, C_{higher}). Though calculated in real time at each time step, energy costs are evaluated in units of cost per hour.

Fan energy cost per hour C_{fan} (\$/hr) = P_{fan} * R_e * 1 hr

Cooling energy cost per hour C_{cool} ($/hr) = P_c * R_e * 1 hr

Where R_e is the utility electricity rate in $/kWh. The R_e rate may vary according to time of day, time of year and/or volume block depending on local utility rates, but only energy charges are included. This approach cannot directly account for electricity demand charges.

Heating energy cost per hour C_{heat} ($/hr) = P_h * (1 therm / 100,000 Btu) * R_g * 1 hr. For electric heating systems, C_{heat} ($/hr) = P_h * (1 Btu / 3412) * R_e * 1 hr

Where R_g is the utility natural gas rate in $/therm. The R_g rate may vary according to time of day, time of year and/or volume block depending on local utility rates.

Total HVAC energy cost per hour C ($/hr) = C_{fan} + C_{cool} + C_{heat}.

Humidity Control (Optional)

If humidity control is applied, during occupied mode and setup mode, the upper limit of the SAT setpoint shall be constrained by outdoor air dew point with the following parameters.

Variable	Value
SP^h_{min}	Min_SAT^h
SP^h_{max}	Max_SAT^h
OAD_{min}	Min_OAD_
OAD_{max}	Max_OAD

If the outdoor air dewpoint (OAD) is lower than OAD_{min}, the SAT shall not be higher than SP^h_{max}; if the OAD is higher than OAD_{max}, the SAT shall not be higher than SP^h_{min}; if the OAD is between OAD_{min} and OAD_{max}, the SAT shall not be higher than SP^h, which can be calculated as follows:

$$SP^h = SP^h{}_{max} - \frac{(SP^h{}_{max} - SP^h{}_{min}) \times (OAD - OAD_{min})}{OAD_{max} - OAD_{max}}$$

An example is given below, with Min_SAT^h , Max_SAT^h , Min_OAD, and Max_OAD being 58°F, 65°F, 55°F, and 60°F, respectively.

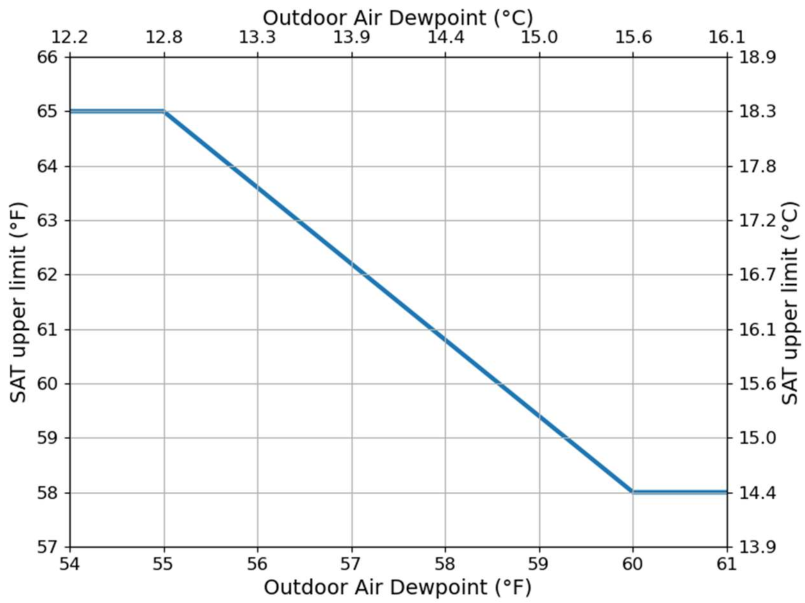

Appendix B: Figures, Tables, and Detailed Information

Figures

Energy model simulation results: supply air temperature, energy, and cost comparison

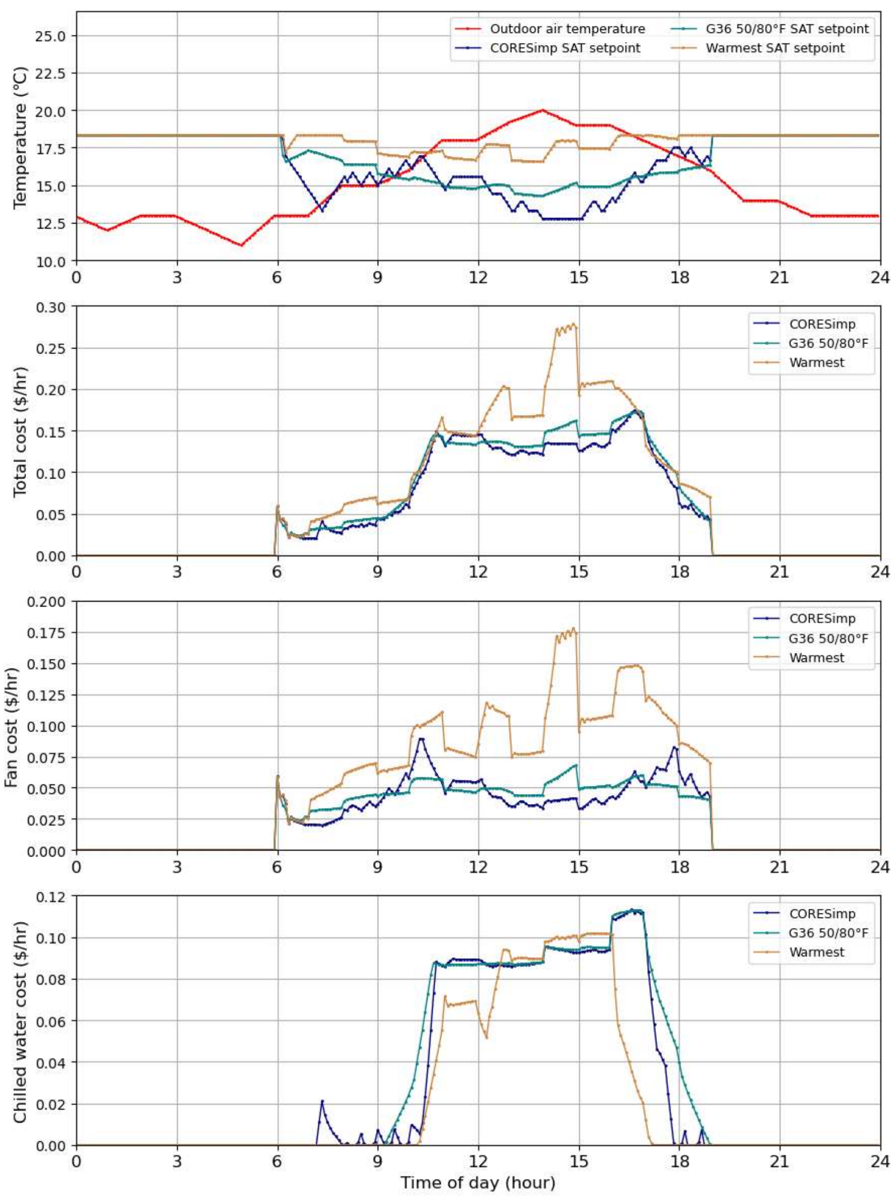

Figure B 1: Snapshots from one June day in the Oakland baseline scenario illustrate how SATs and the total, fan, and chilled water costs varied among CORE 2, ASHRAE G36 with 50/80°F OAT limits, and Warmest

scenarios. The heat energy cost was not shown due to no heating demand on the given day.

Energy and Cost Savings Summary—MOB-1

Following are six violin plots showing the distribution of daily energy use for Baseline and CORE measurements for each RTU.

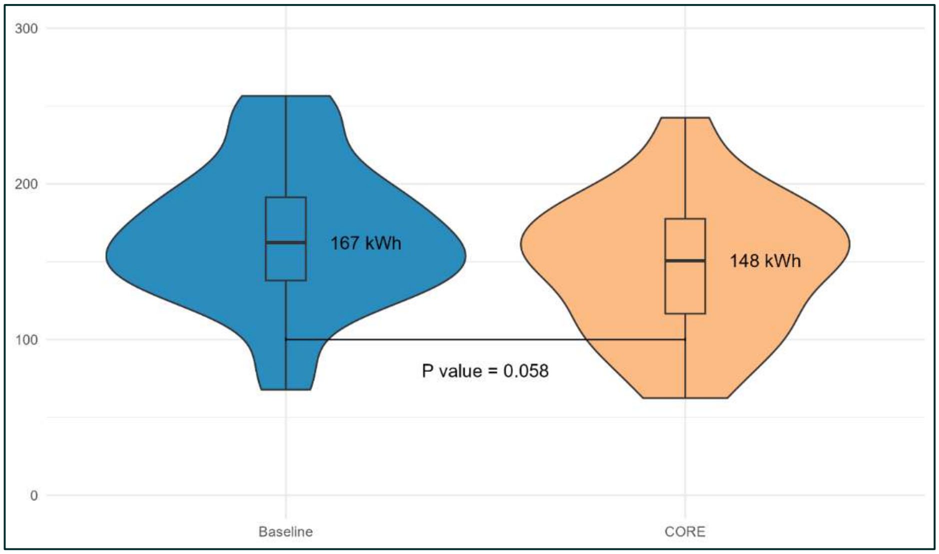

Figure B 2: Violin plot comparing daily energy consumption distributions for Baseline and CORE for RTU 1 – MOB-1.

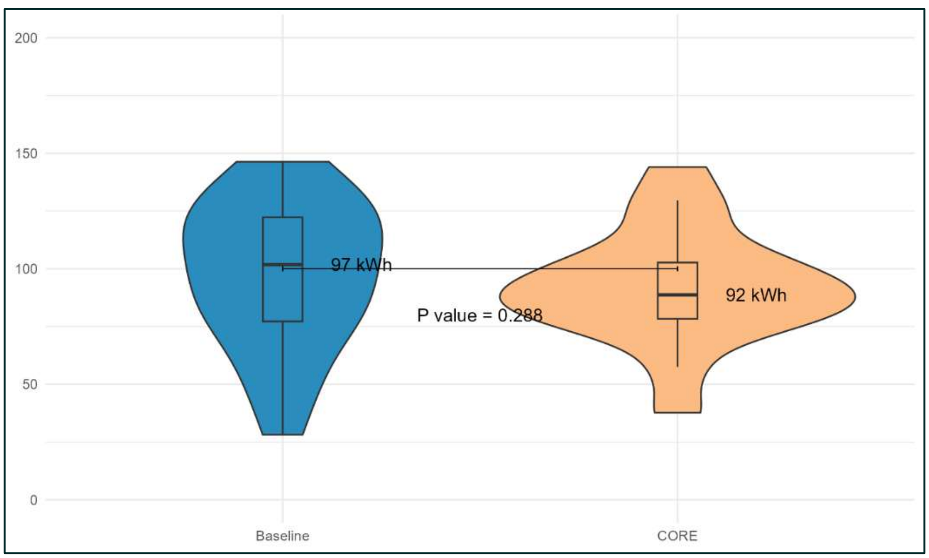

Figure B 3: Violin plot comparing daily energy consumption distributions for Baseline and CORE for RTU 2 – MOB-1.

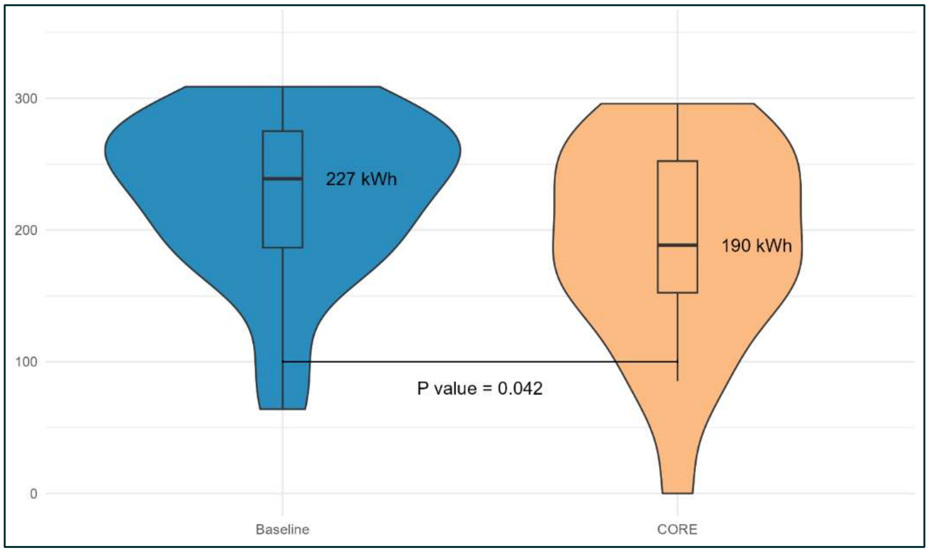

Figure B 4: Violin plot comparing daily energy consumption distributions for Baseline and CORE for RTU 3 – MOB-1.

Violin plots of the distribution of daily energy use for Baseline and CORE measurements for each RTU

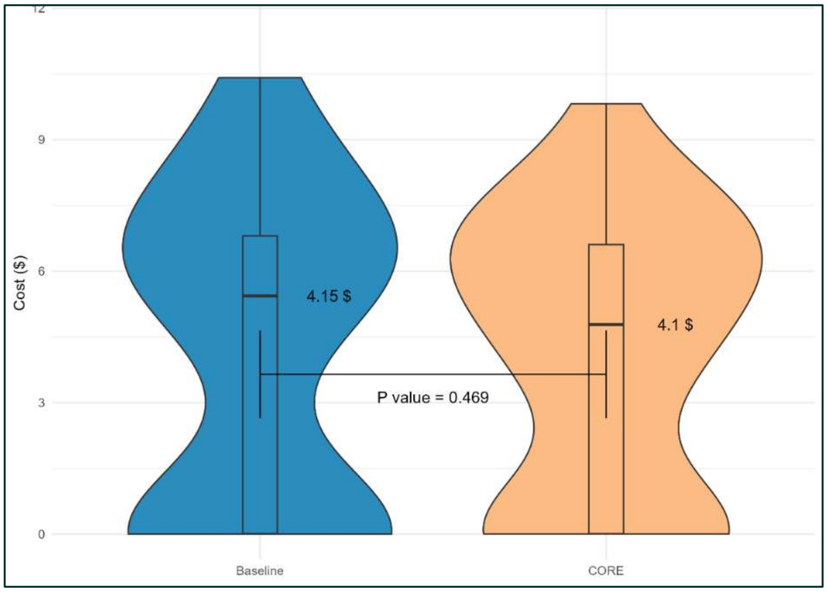

Figure B 5: Daily Heating, Cooling, and Fan energy cost distribution for Baseline and CORE for RTU 1 – MOB-1.

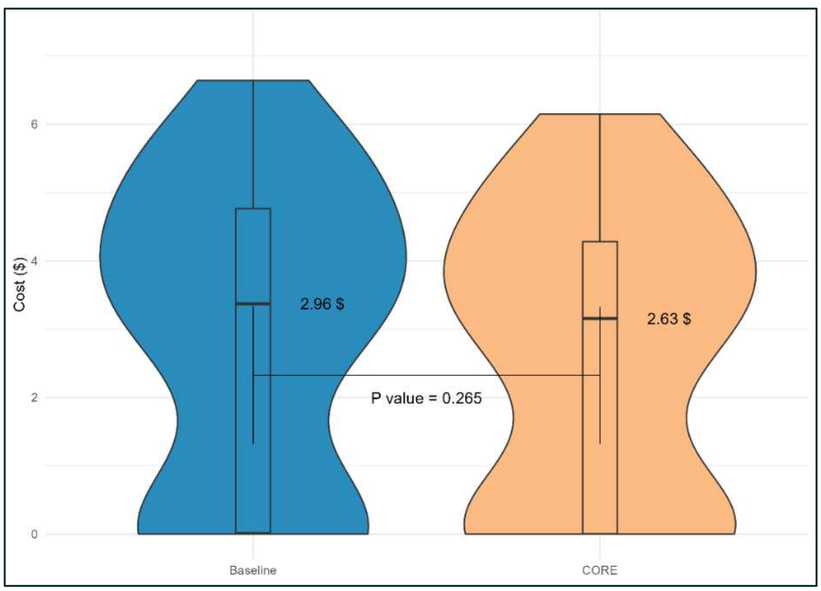

Figure B 6: Daily Heating, Cooling, and Fan energy cost distribution for Baseline and CORE for RTU 2 – MOB-1.

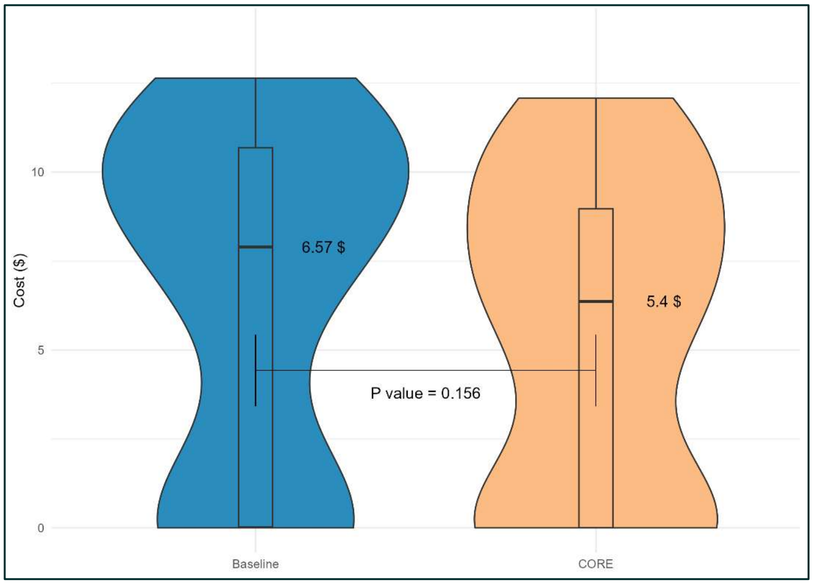

Figure B 7: Daily Heating, Cooling, and Fan energy cost distribution for Baseline and CORE for RTU 3 – MOB-1.

Additional information

From: Core implementation efforts and issues

EXAMPLE OF CALCULATION COMPLEXITY

An example of calculation complexity is with the ΔTc calculation interval, which is intended to represent the long-term average temperature difference between the two closest temperature sensors immediately upstream and downstream of a coil, when that coil is closed and has been closed for long enough for any fluid in the coil to reach ambient conditions in the duct. The team expected the value of ΔTc to be relatively stable. During commissioning, the team observed that the value of ΔTc changed rapidly to values that were 10 times what was expected, causing unusually high cooling power values, leading the algorithm to incorrectly move towards warmer SAT setpoints. While the team observed the high ΔTc values, a significant amount of time was spent troubleshooting before identifying the cause. Calculation equations, inputs, and the constant value were all confirmed. The calculation at a snapshot of time was verified as part of the functional performance testing. A manual calculation was performed using downloaded trend data and, by further investigating, the team found a discrepancy in the updating interval of ΔTc. While the CORE algorithm was calculating and returning new SAT setpoints at five-minute intervals, the ΔTc calculation was occurring every 2.5 seconds, leading to multiple updates to ΔTc between each SAT setpoint calculation using the same parameters for the equation. This frequency led the ΔTc value to escalate faster to large values. Changing the ΔTc calculation interval to five minutes fixed this issue.

ΔTc value in the cooling power calculation is a temperature correction to account for fan heat, sensor drift, and/or passing control valves. As an example, if the fan motor releases heat to the supply air, the cooling coil would have provided more cooling than what would be calculated based on the supply air temperature and the temperature difference caused by this fan motor heat gain would remain relatively stable. ΔTc is calculated using an exponential average equal to [k * (Ts − Tm) + (1 − k) * (ΔTc from last time step)] calculated during periods when the chilled water control valve has been closed for a minimum of 5 minutes and when airflow is proven. The value of ΔTc is fixed at its last value prior to the valve opening and for 5 minutes after closing. The exponential smoothing coefficient k is user-adjustable, with a default value between 0.01 and 0.001.

EXAMPLE OF AN EXISTING HVAC SYSTEM, BUILDING DESIGN, OR CONTROL ISSUES

Another example is MOB-2 RTU-2, which had a non-functioning economizer due to an OAT lockout restriction on cooling coil operation that prevented the ability to use free cooling from economizing. This meant that the SAT could not meet the SAT setpoint for the combined RTU 2/3 unit, and that the zones received warmer air and more frequently requested cooling. These factors led the algorithm to move out of the cost optimization path. The team expects the correlation between energy use and temperature to be poor and the savings to be diminished.

It is not fixable as part of this effort, and the team expects this to greatly decrease the cost savings achievable by CORE since additional energy is used to reheat the supply air.

ADDITIONAL DETAILS ABOUT THE CHALLENGES OF WORKING WITH AN EXISTING CONTROL SYSTEM

A detailed review using manual calculations found that the Pc calculation using the equation Pc = Pchw * Ec * (1/12000) was not working correctly in the BAS. We could not determine if this was a bug in the program or an order-of-operations issue, but the equation Pc = (Pchw * Ec) / 12000 produced correct values. During the monitoring period, Pc values in MOB-1 showed zero for periods of time for each RTU due to a non-value being assigned. This caused these periods to not run CORE correctly. The team added a workaround to use the last known value if a non-value is passed.

AN ADDITIONAL EXAMPLE OF ISSUES OF WORKING WITH AN EXISTING CONTROL SYSTEM

Another example of this issue was highlighted in reviewing the number of request ignores in MOB-1 and MOB-3. At MOB-1 and MOB-3, the team found that the cooling coil turned on more frequently than expected and that SAT setpoint did not move towards reducing cost. Further investigation showed that CORE was using an incorrect, lower number of ignores than specified (same number of ignores as the baseline). This led to the algorithm coming out of the cost optimization path (right side path of Figure 1) greatly reducing the potential cost reductions. The team found this difficult to investigate due to multiple variables being mapped incorrectly previously at the site; and resolved the issue by using new variables.

Request ignores are used in ASHRAE G36 trim and respond logic where the cooling or static pressure requests need to be more than the specified ignores limit for the system to respond to cooling/pressure. The purpose of this is to avoid a few rogue zones requesting cooling or airflow all the time. ASHRAE G36 recommends 10 percent of the total requests as number of ignores.