Title
Cloud-Control of Legacy Building Automation System: A case study

Authors
Prakash, Anand Krishnan
Pritoni, Marco
Kloss, Margarita
et al.

Publication Date
2022-08-01

DOI
10.20357/B7BS4H

Peer reviewed

Cloud-Control of Legacy Building Automation System: A case study

Anand Krishnan Prakash[1], Marco Pritoni[1], Margarita Kloss[1], Mary Ann Piette[1], Michel Kamel[2], Dotty Hage[2]

Lawrence Berkeley National Laboratory
Melrok LLC

Energy Technologies Area
August 2022

[Insert citation or DOI # here]

This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy,
Building Technologies Office, of the US Department of Energy
under Contract No. DE-AC02-05CH11231.

Cloud-Control of Legacy Building Automation System: A case study

Anand Krishnan Prakash, Marco Pritoni, Margarita Kloss, Mary Ann Piette, Lawrence Berkeley National Laboratory
Michel Kamel, Dotty Hage, Melrok LLC

ABSTRACT

As Internet of Things devices and cloud-based platforms become more mature, Energy Management and Information Systems (EMIS) are increasingly gaining momentum in the building industry. In large commercial buildings, Fault-Detection and Diagnostic (FDD) and energy information systems (EIS) are now established technologies with tens of providers and thousands of deployment sites across North America. The new frontier for the EMIS technology is now represented by control systems that use advanced system optimization (ASO) methods to improve the operations of the HVAC system. Given the complexity of the integration of such systems with the existing building automation systems (BAS) and the higher risk involved with direct control of the HVAC, these systems are still emerging in the market.

This paper presents the results of a project in which a start-up company partnered with a research institution to develop a cloud-based software EMIS solution and deployed it in a university campus in California. The software system included advanced sensing, data acquisition, storage and advanced control and analytics applications developed on top of the native BAS. The new platform controls ten buildings on the campus and the FDD and the ASO applications deployed on this platform were able to generate energy savings of up to 35% and 25% in certain buildings for each functionality respectively. Where the platform did not save energy, it improved building service (air quality). Lessons learned include the importance of collaborating with and training the building operators and evaluating whether the legacy system can work reliably with the new technology.

Introduction

In a study of 24 university buildings in California (Mills 2009), buildings were observed to “drift” from their optimum performance achievable through continuous commissioning. This drift can lead to more than 40% waste in the building’s total energy consumption. Even so, the return on investment in costly manual interventions to correct drift is not compelling for most building owners, given that re-commissioning may need to be repeated as often as every few years to sustain the savings. Unfortunately, conventional Building Automation Systems (BAS) fail to optimize energy use because predetermined settings become rapidly obsolete, and most BAS are not able to continuously and automatically optimize set points for key systems, such as variable frequency drives, valve positions, and damper positions. Conventional BAS also do not respond dynamically to changes in building schedules and room occupancy, and they are not cognizant of current and forecasted environmental factors and grid conditions that would enable optimization of energy use.

To address these challenges, several Energy Management and Information Systems (EMIS) solutions have been developed and are now available on the market for the building industry. Most of the EMIS solutions that exist today serve as a data acquisition, storage and analytics system that is user friendly and that is more accessible to more than just the energy engineers on-site (Kramer 2020). Through the collected data, these platforms are able to provide the customer (building owner, occupant, facility manager etc.) with predictions of building energy consumption, future costs and also suggest improvements and the corresponding potential benefits. However, we are now seeing significant advances made on these EMIS systems to provide more than just analysis and visualization capabilities. Latest research literature and commercial products have started to control the BAS directly by overriding some of the programmed setpoints themselves (O'Grady 2021, Pritoni et al. 2022). Through this, they could be able to continuously optimize the management of building sensors, energy consuming devices, and existing energy management systems through advanced analytics and control capabilities.

Existing literature (O'Grady 2021) shows that the market penetration of such technologies is still in its infancy. Hence, to further the research and commercial adoption of such platforms, we present a case study where we developed and deployed such a middleware EMIS platform in a university campus in southern California, continuously optimizing the operation of multiple buildings. In this paper, we will focus on ten representative buildings out of all the buildings on the campus and will present a snapshot of the software applications that were deployed, the benefits reaped and the challenges encountered. The main contributions are:

- the technology stack that can be deployed across a whole campus to collect data and run Fault Detection and Diagnostics (FDD) and Advanced System Optimization (ASO) applications in real-time
- quantitative benefits of deploying such a solution in a campus
- a discussion of the challenges and major roadblocks encountered, along with possible solutions. This could be transferable to another deployment site

The rest of the paper is organized as follows: we introduce the campus and the different buildings of interest, followed by the technology stack that was deployed - both hardware and software. The applications that were demonstrated and the evaluation of these applications are the next two sections. Then we present a discussion of the key challenges and takeaways from this demonstration effort, followed by the conclusions.

Site Description

The ten buildings on a university campus in southern California that were part of this deployment had a total area of approximately 280,000ft^2. They included various non-academic buildings: one administrative building with office spaces (B1), a data center (B2), a student activity center (B3), a music hall (B4), a library (B5) and five academic buildings containing classrooms and work spaces (B6, B7, B8, B9 and B10). B4 is an all-electric building that is being conditioned all the time to maintain the right environment for the musical equipment set up in this building. All other buildings were conditioned by district heating (gas-based boilers) and cooling (chillers). These ten buildings together had an average energy use intensity (both electricity and gas) of 159kBTU/ft^2 and contributed 23.6% to the campus's annual energy consumption in 2018.

Most of these buildings had a BAS using the standard BACnet/IP (Building Automation and Control Network over Internet Protocol) protocol to control the heating, ventilation and air conditioning (HVAC) operations. However, a few of them used a BAS that communicated using

Schneider Electric's proprietary Infinet protocol. All the power meters and the energy meters that were installed communicated using IP-based Modbus TCP/IP (Transmission Control Protocol over Internet Protocol) or the serial Modbus RTU (Remote Terminal Unit) protocols. The buildings had building-level meters and submeters to measure individual equipment (chiller, boiler etc.) power consumption. The BAS controlled central chillers and boilers that supplied chilled water and cold water to groups of buildings, Air Handling Units (AHUs) and Variable Air Volume (VAV) boxes. In some buildings (e.g., laboratories), the BAS also controlled fume hoods to exhaust air from the spaces. In these ten buildings, a total of 339 BAS controllers continuously monitor and control almost 4,000 (measurement and control) data points. All buildings on campus also had a Wireless Local Area Network (WLAN) set up using Cisco WLAN infrastructure. It manages more than 1000 access points (AP) for the occupants (nearly 1700 students, faculty and staff) to access the local intranet and Internet.

Technology Stack

This section describes the software and hardware infrastructure that was developed and deployed in the buildings at the university campus. Figure 1 provides an overview of this platform, both software and hardware.

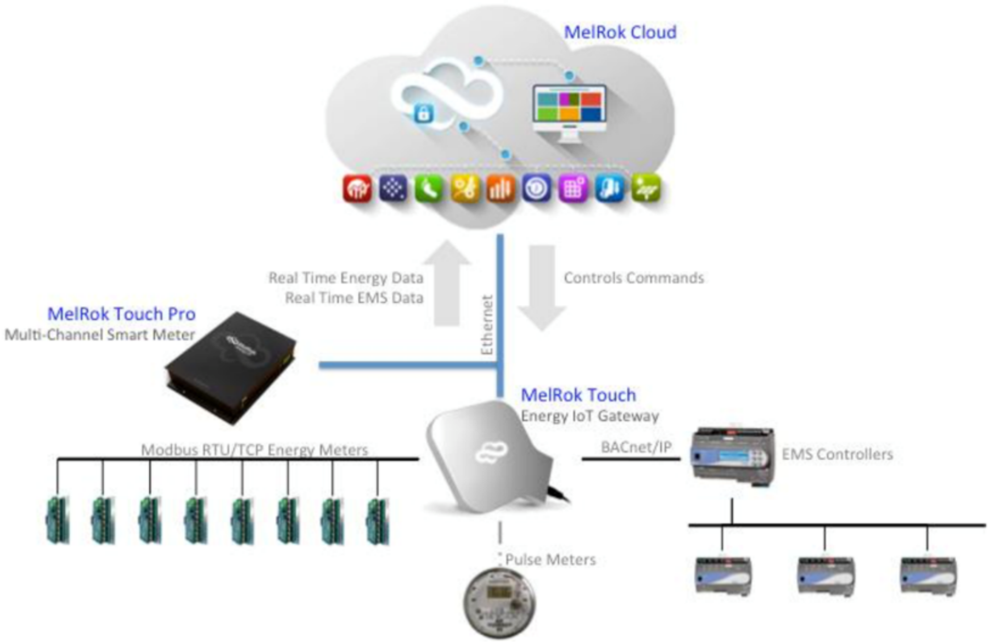

Figure 1. Overview of the Melrok technology stack that was developed and deployed at the university campus.

Local Gateways

We installed a local gateway, Melrok Touch (Touch gateway) (MelRok Energy IoT 2019) in each building to connect with the on-site BAS controllers and energy meters. Melrok Touch gateway has the capability to communicate across different network protocols such as Modbus TCP, Modbus RTU over RS485 and BACnet/IP and this enables bidirectional communication with these devices.

Cloud Infrastructure

The energy readings and the BAS data collected by the local Touch gateways were uploaded to the Melrok's cloud infrastructure once every minute. The data was stored in an Apache Cassandra database and we used Apache Kafka to trigger real time analytics, FDD and ASO applications. In case of loss of Internet connectivity, the Touch gateways temporarily stored the data collected locally before reattempting to push it to the cloud database once the network is restored.

Occupancy Counting Using Anonymized Wi-Fi Data

One of the key factors of energy consumption in commercial office and academic buildings is the occupancy count (Chen 2018). This data can be used for several applications such as developing energy models to predict future energy consumption and preventing energy wastage when rooms are being conditioned even though they are unoccupied. However, as there were no occupancy sensors previously installed in any of the buildings at the university campus and it was expensive to install new ones and integrate it with the BAS, we used the number of devices connected to the Wi-Fi network in a building as an approximation of the total building occupancy count. We used the Counting Occupants Using Network Technology (COUNT) (Clark 2020) open-source software for acquiring this data. COUNT runs SNMP[1] queries on the WLAN controller to obtain the number of devices connected to each access point at any instant. Then, it anonymizes and aggregates the data to obtain a total building occupancy count.

It is to be noted that Wi-Fi based occupancy count could lead to over-/under-counting of the actual occupants. Our previous work (Clark 2020) had looked into this and saw that overall, the Wi-Fi trends correlate well with other indicators of occupancy such as class schedules, CO2 sensors, and manual counts of people. Therefore, for buildings without occupancy counters, this approximate count provides an occupancy profile that closely matches the actual number of occupants in the building.

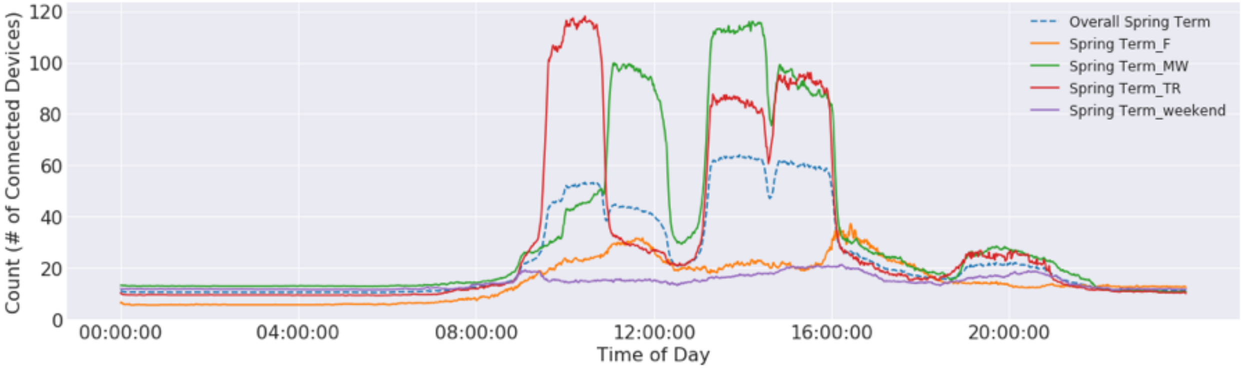

Figure 2. Average occupancy patterns during different days (MW - Monday, Wednesday, TR - Tuesday, Thursday, F - Friday) in the spring term for an academic building on campus. (Clark 2020)

---

[1] Simple Network Management Protocol (SNMP)

Data Acquisition

The installed Touch gateways were used to query the BAS and the energy meters (once every minute) and pushed the retrieved data to the Melrok Cloud. We also set up COUNT to obtain the approximate occupancy count in all the buildings and pushed it to the Melrok cloud database. Figure 2 shows an example of this approximate occupant count for one academic building on campus during the spring semester of 2019.

One of the main challenges in setting up this infrastructure was integrating with BAS controllers that communicated using Schneider Electric's proprietary Infinet protocol. To work around this limitation, we used secondary BACnet controllers as gateways to mirror the Infinet controllers. Through this configuration, the Touch gateway queried these intermediate BACnet controllers using BACnet/IP to acquire the values from the BASand to write control points back to them.

The status of all connected devices was monitored by the MelRok cloud dashboard, with email notifications pushed out in case of lost connectivity. This approach ensured that all the meters were kept online and, in case of any downtime, work orders were generated and the issues resolved.

Assignment of Standardized Metadata to Points

Another major challenge encountered during the project was interpreting the meaning of the measurement and control points in the BAS, given the lack of consistency in the naming convention. This is actually a fairly common issue, recognized in the literature (Pritoni 2021) as a major challenge obstructing the adoption of scalable and portable data analytics and controls in buildings. While a few solutions were emerging to address this lack of standardization (e.g., Project Haystack (Project Haystack 2022), Brick (Balaji et al. 2016)) and an ASHRAE standard was under development (i.e., ASHRAE 223p: ASHRAE 2018), at the time of the project these initiatives were inadequate or incomplete. Therefore, for practical reasons, we adopted an internally consistent naming convention across all buildings, mapping existing point names in the BAS to this schema. All the BAS points continuously read and written from the BAS were associated with a consistent type (e.g., supply air temperature sensor) and mapped to a global unique identifier (GUID). With this approach an application can request all points of a particular type and use the GUIDs returned to then query the time-series data related to that point. Without such mapping, it would have been significantly more difficult to automate analytics. The mapping process was done partially automatically, for 75% of the BACnet points and 40% of the Infinet points, and manually for the remaining ones.

Applications

This section describes the different applications that were developed and deployed in Melrok cloud, which in turn provided feedback in the form of analytics and metrics to the user dashboard, or control signals to a specific BAS equipment in a building. Results from these applications tested are provided further below.

Visualization and Analytics

The Melrok cloud platform provided a web-based interface to query and analyze the data retrieved from the local BAS controllers. Data can be viewed as tables and as time-series, with a host of visualization tools such as heat maps, box plots, profile plots, scatter plots, and load duration curves. Results can be easily filtered, sorted and downloaded. Energy managers and the facilities staff at the university were given access to this dashboard, to analyze the statuses and performance of the different control systems in the campus. Figure 3 is one such visualization that shows a day's trend of the cooling coil valve positions for all AHUs in a building.

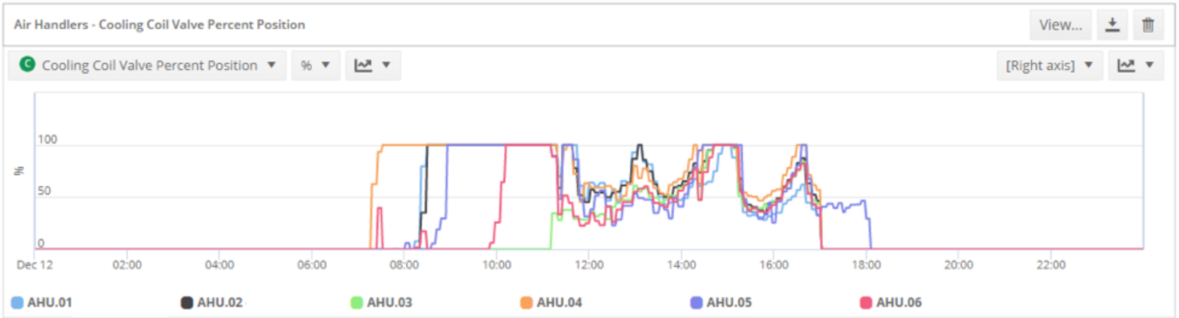

Figure 3. The cooling coil valve positions of the different AHUs in an academic building (B7) throughout a day

The cloud platform also included a python Jupyter sandbox (Kluyver 2016) for data querying and exploration. We used this environment to develop several applications such as occupancy aware dynamic HVAC schedules and load forecasting. The occupancy aware scheduling has the potential to reduce HVAC use as energy managers can identify when the HVAC system is running, but nobody is present in a building. For example, Figure 4, shows that the Thanksgiving break was not programmed as a holiday and the building operated as usual (indicated by the AHU1 and AHU2 Supply Fan statuses) even when the campus was closed (Clark 2020). Additionally, identifying these anomalies in real time enables managers to address them without having to program them into the BAS schedule, which is a manual and opaque process.

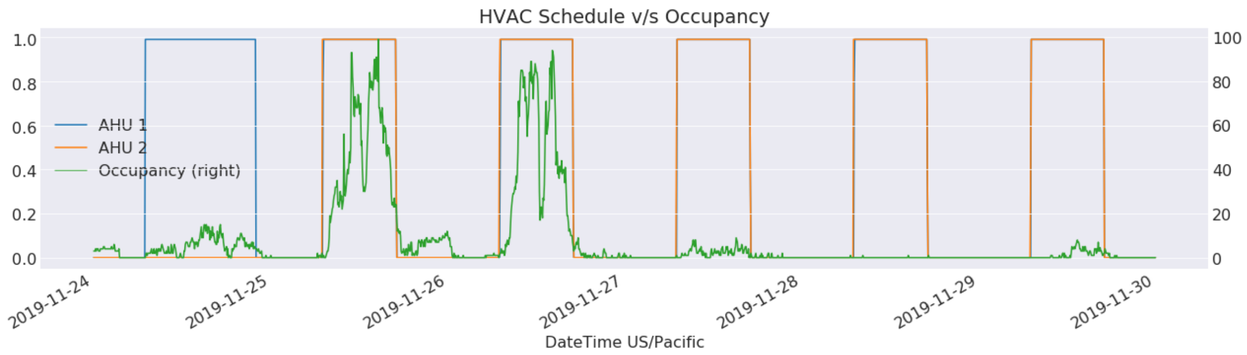

Figure 4. Occupancy and HVAC schedule at a building before (Sunday-Tuesday) and during the Thanksgiving break (Wednesday-Friday). Left axis: AHU supply fan status (0=off, 1=on); Right axis: Occupant count (Clark 2020).

On-Demand Commissioning

A set of auto-commissioning modules was also included in the cloud platform. These analyze BAS data to detect problems and inefficiencies. They sift through the one-minute data from all systems in a building to detect feedback that does not match control commands, metrics that are not set back at night, equipment that have no reset strategies and terminal units that supply excessively hot air (in violation of Title 24 (California Energy Commission 2018)) and then notify the energy manager as needed.

Fault Detection and Diagnostics

We were able to implement FDD through simple rule-based identification and data analysis. Rule based FDD allowed us to detect and diagnose sensor faults such as: frozen sensors (unchanging data for a long duration), missing data (non-reporting sensors), out-of-range data (function of metric type and units), and short cycling sensors. The most common sensor failures were frozen sensors, followed by short cycling sensors. Missing data was often the result of network outages. Two Touch gateways also malfunctioned due to power-outage related voltage spikes. When a Touch gateway is down, the points being acquired by the Touch are re-routed to another Touch gateway.

Data-driven fault detection algorithms consisted of seven categories of failures: simultaneous heating and cooling, inconsistent temperatures, failure to heat, failure to cool, defective dampers and ventilation failures. Results of the fault detection were displayed on the cloud dashboard, with bar charts highlighting the times the failure was detected. Timeline charts of the relevant metrics were also displayed on the same dashboard page for quick diagnosis of the problem. The fault detection algorithms were calculated in real time at one-minute intervals. When a problem was detected, it was first flagged as a violation. Violations that last 15 minutes are tagged as faults. Faults that last 2 hours or more were then tagged as Alerts and notified to the energy manager. Once the energy manager received the notification, they would then either update or fix the BAS's sequence of operation or initiate a work order in scenarios where parts had to be replaced.

Advanced System Optimization

The ASO applications that we developed were deployed in the cloud platform. They continually analyzed the operational data and generated control signals that were sent to the local BAS controllers through the corresponding Touch gateways. These applications overrode the sequence of operations that were programmed on the controllers. One such application was the dynamic optimization of air handlers in a building, specifically the outside air damper position, the supply air temperature and the supply air static pressure. The standard control sequence utilized by the legacy BAS did not implement any resets for these setpoints. The new AHU optimization used real-time data from the air handlers and from all terminal units being served by the particular AHU, including the heating and cooling demand for each zone. The application dynamically changed these setpoints to holistically optimize the electricity use for fans and chillers with constraints on thermal comfort (i.e., zone temperature) and air quality (i.e., outdoor air fraction to the zone).

Evaluation

Measurement and Verification Plan

The evaluation of the savings and performance of the platform was more difficult than expected. Even though we had developed a measurement and verification (M&V) plan at the outset of the project (for both the academic and non-academic buildings), at the end of the implementation period, we had to alter the proposed evaluation approach. Changes in building operation and occupancy due to COVID-19, made the baseline building energy use, collected before the deployment of our solution, not directly comparable with the energy use in the post-implementation period. For most facilities in the university campus, the occupancy and schedules were substantially altered under COVID-19 lockdown since instruction was moved online. To overcome this obstacle, we developed three methods to evaluate savings:

- Method 1 (M1): A multi-variable empirical baseline model was created for each building, where model features included hourly temperature, time of day and day of week. The model was trained using data during the year 2018. This method was used to effectively measure the impact of the commissioning efforts undertaken as a result of the FDD applications that were developed. This includes replacing faulty and erroneous sensors, updating few equipment and certain sequence of operations in the BAS controllers.
- Method 2 (M2): After a few months in the lockdown, while most of the academic buildings remained unoccupied and without any conditioning, staff returned to a few buildings. This allowed us to collect data for a short period of time of baseline operation. This data was not very representative because improvement measures guided by FDD applications had already been implemented in some buildings. Nevertheless, it still provided a measure of performance with no on-going cloud-based ASO.
- Method 3 (M3): This method consisted of cycling the ASO application ON and OFF sequentially for short periods (e.g. 1 day) to observe any difference in building (or equipment) behavior with and without the cloud-based control optimization. On its own, this method cannot yield results that can be extrapolated to annual savings, since it can only be used to compare two periods with similar levels of occupancy and weather. However. it can be used to spot check changes in a building's behavior (or in the operation of specific equipment) due to the difference in control logic between BAS controls and cloud based ASO control.

Fault Detection and Diagnostics Results

A total of almost 260 formal work orders had been issued since the platform was deployed. Of these, some 57 work orders were still open as of April 2021, when this project was completed. COVID shutdowns reduced both the flow of work orders and their completion. Figure 6 is a breakdown by category of all closed and open work orders. Malfunctioning actuators make up the majority (55%) of work orders followed by faulty sensors.

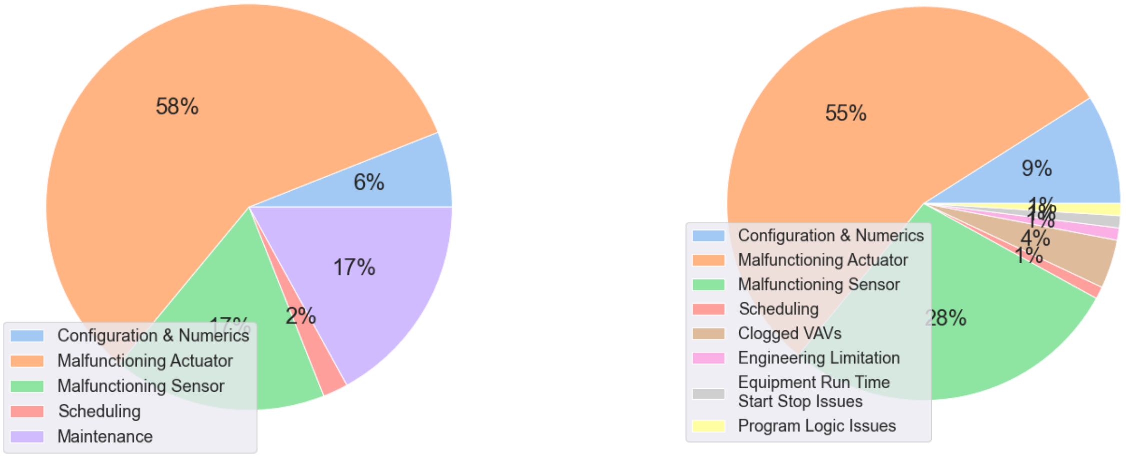

(a) (b)

Figure 6. Categories of different (a) closed and (b) open work orders that were issued due to the FDD alerts

The access to 1-minute BAS data also allowed the platform to detect the short cycling of equipment (continuously turning on and off, which dramatically shortens the chiller life cycle) and to notify the facility staff. One instance of this was seen in the new 120-ton chiller at the data center building (B2). Once this anomaly was flagged (Figure 7), it was turned in as a work order via the university facilities system, with the work order executed on March 3rd 2020. A review of the trends at 1-minute resolution revealed that the chiller had been short cycling since its installation in fiscal year 2016.

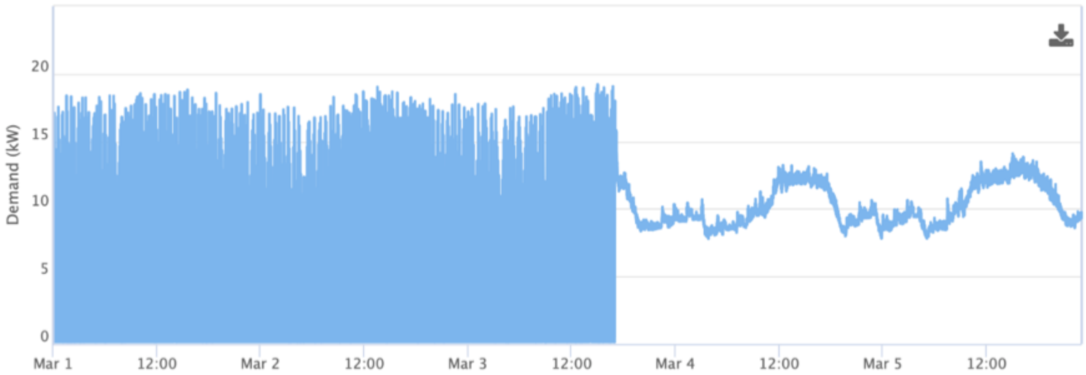

Figure 7. Power profile of a short-cycling 120-ton chiller before and after this issue was fixed on March 3.

Advanced System Optimization Results

Figure 8 illustrates the performance of the cloud based advanced control strategies for the two AHUs in an academic building by comparing their operations on February 28, 2020 (without ASO) to March 13, 2020 (with ASO). Figure 8(a) displays the line plots for supply air static pressure on when the ASO was off (dotted blue and black lines) and after it was turned on (solid blue and black lines). The drop of about 0.5 inWC in supply air static pressure was accompanied by a decrease in supply air fan current from about 6A (dotted lines) to 4A (solid lines), or 30%, as shown in Figure 8(b). Similarly, the supply air temperature was on average higher during the operation hours with ASO, as shown in Figure 8(c), while all zones were still within their desired space temperature set points.

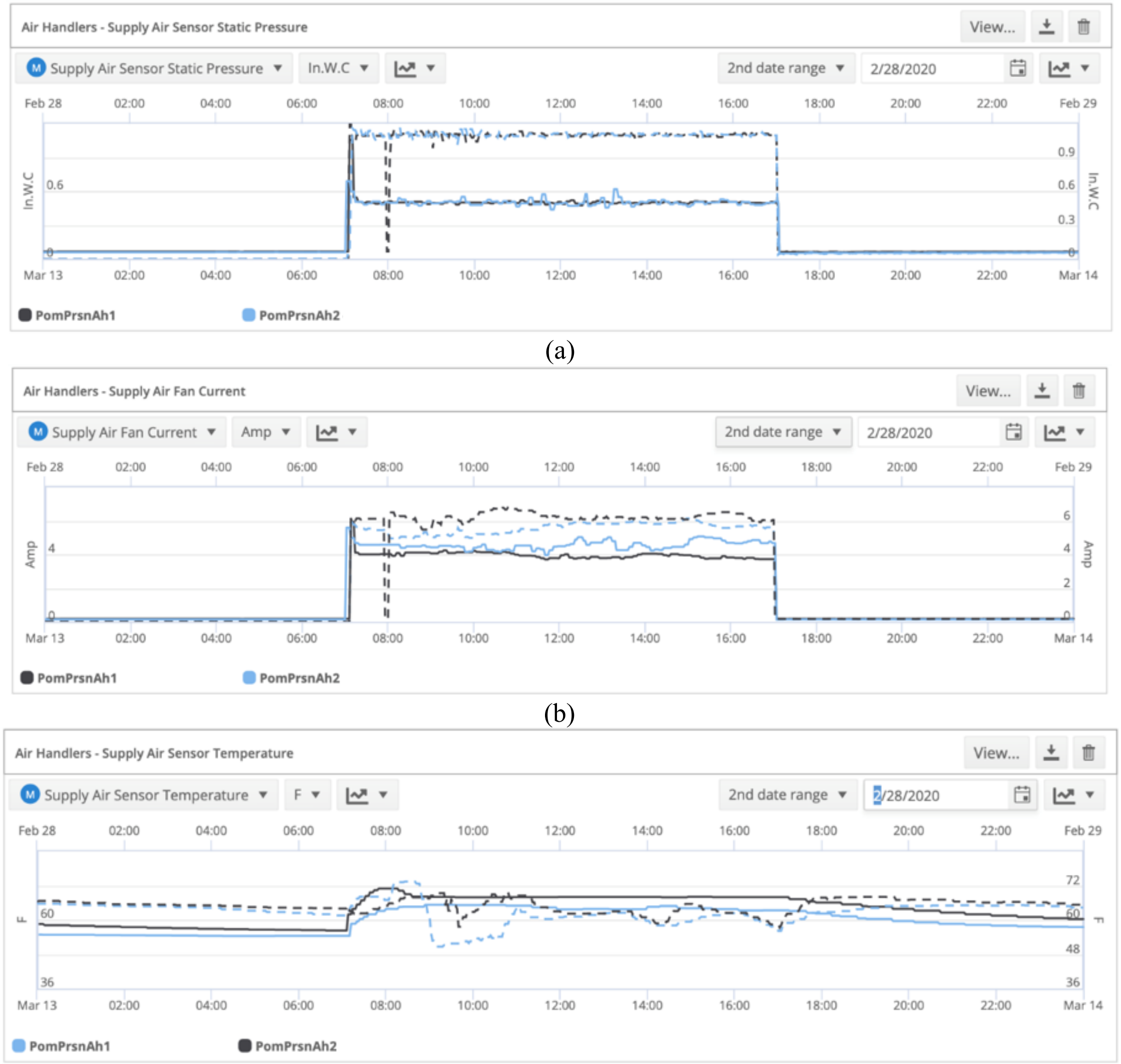

(a)

(b)

(c)

Figure 8: Impact of cloud-based optimizations on the AHUs in an academic building (B9) indicated by (a) drop in supply air pressure (b) drop in supply air fan current and (c) an average increase in supply air temperatures while maintaining zone temperatures within the deadband. The dashed lines represent the baseline scenario and the solid lines represent the optimized control scenario.

Table 1 and Table 2 present a summary of the estimated change in the energy consumption in the ten buildings of interest through the FDD and ASO applications. It also includes the M&V method used for this evaluation (M1, M2 or M3) and the time periods used for calculating baseline and the system performance. From these tables, it can be seen that the platform enabled savings in two ways: 1) by providing information about faulty and underperforming equipment to the campus commissioning team, resulting in savings up to 30% and 2) by directly controlling and optimizing the HVAC operation, providing savings up to 25%. We consider these encouraging results as the best estimate of the savings for this project, but we acknowledge that the methodology for assessing them is not as robust as originally planned and that changes in the baseline related to COVID-19 are hard to assess.

Table 1. Summary of savings due to FDD and ASO in the non-academic buildings: the administrative building (B1), the data center (B2), the student activity center (B3), music hall (B4) and library (B5)

Application	Baseline period	Performance period	Method	Difference in energy consumed (%)				
				B1	B2	B3	B4	B5
FDD	07/17-06/18	07/18- 06/19	M1	-3	14	-15	32	-10
ASO	03/20-07/20 (when occupied, ASO off)	03/20- 07/20 (when occupied, ASO on)	M2	17	N/A	17	N/A	-13
ASO	08/20-12/20 (when occupied, ASO off)	08/20- 12/20 (when occupied, ASO on)	M2	-11	N/A	12	N/A	25
ASO	08/20-12/20 (ASO off)	08/20- 12/20 (ASO on)	M3	15	N/A	10	N/A	20

We believe that the calculated savings for FDD (using M1) represent the lower end of the range of the savings because not all the commissioning efforts as suggested by the FDD applications had been implemented yet, particularly due to the delay in executing repairs in certain buildings (as can be seen from the Figure 6(b)).

Table 2. Summary of savings due to FDD and ASO in the academic buildings (B6-B10)

Application	Baseline period	Performance period	Method	Difference in energy consumed (%)				
				B6	B7	B8	B9	B10
FDD	07/17-06/18	07/18- 06/19	M1	1	-2	9	-4	8
ASO	03/20-07/20 (when occupied,	03/20- 07/20 (when occupied, ASO on)	M2	N/A	N/A	N/A	N/A	N/A

	ASO off)							
ASO	08/20- 12/20 (when occupied, ASO off)	08/20- 12/20 (when occupied, ASO on)	M2	N/A	19	-1	7	N/A
ASO	08/20- 12/20 (ASO off)	08/20- 12/20 (ASO on)	M3	N/A	25	20	13	N/A

Lessons Learned

While the technology stack we developed and used to evaluate the benefit of the EMIS was able to produce savings in most buildings, there were significant obstacles and delays that we encountered which slowed down progress. We have encountered some of these barriers in other demonstration projects as well and hence we hope that documenting them in this paper will provide future research projects and commercial solutions in this domain of advanced control platforms an idea of what to expect and how to navigate them.

- **Interfacing with legacy systems and proprietary protocols:** Oftentimes legacy systems use proprietary protocols with little documentation or customer support. They may use serial communication and hence are much more susceptible to network reliability issues and often result in data loss. Hence, it is important to constantly monitor the status of these points to quickly detect any failure in communication. For new BAS installations and upgrades, the use of IP-based and published protocol systems is critical for ease of interoperability.
- **Assigning semantic information to BAS points:** It is important not to assume that the point names accurately reflect what they are monitoring or controlling. There were several examples of points that were mislabeled in the BAS and the mislabeling was never detected. While the tagging to a standard naming convention accurately reflected the point's name, the data was not consistent with expectations. An example was the mislabeling of chilled water inlet and outlet in the BAS.
- **Costs:** The main expenditure in deploying such an EMIS solution are due to the hardware purchases, the commissioning process and the recurring cloud service fee. The recommended one gateway per building can cost around $500 (depending on the number points to be queried) and another $1000 for commissioning this gateway. However, in this case study, setting up interfaces to some of the proprietary systems required more effort and hence, incurred higher charges. Based on the set of applications that have been deployed, the software fee can vary between $50-$150 per month. Note that as the commissioning becomes more streamlined, the costs and effort will start to reduce.
- **Cybersecurity and privacy concerns:** Most campuses require the BAS controllers and the HVAC equipment to be on a separate network than the normal campus LAN network for cybersecurity concerns. Hence a new solution provider could face resistance when

they request access to this network to install local gateways. Privacy concerns also surfaced when we tried to install occupant counters in a building. This experience points to the need for future cross-training between energy and IT staff as building systems become more software-driven and could-integrated.

- **Surfacing of Hidden Problems:** While connectivity to all BAS points is critical for cloud-based optimization, it highlights a number of existing failures that have previously gone unnoticed. Some of these failures are logical, such as the mislabeling of points discussed above or when a control program freezes randomly. Other failures are physical in nature, with broken communication wires, defective sensors, defective actuators, etc. We recommend budgeting for the time to fix these issues in advance so that there are no surprises.
- **Vulnerability to Device Failures:** It is always recommended to have a fallback mechanism in the local gateways to return the control back to the local BAS controllers whenever critical faults that may impact the cloud-based control (such as defective sensors, missing data, loss of network connectivity etc.) are encountered.
- **Start controls as soon as possible:** After the FDD application generated a large number of work orders, we waited for the retro-commissioning efforts to conclude before deploying our ASO applications. This resulted in delaying the commissioning of cloud-based controls in almost half the buildings, even though the benefits of these controls were obvious, even in buildings that were not retro-commissioned. A lesson learned for future implementations is to proceed with cloud control at the earliest opportunity to ensure that we obtain the benefits from low hanging fruits such as dynamic schedules and improved control algorithms. As cloud control proceeds, faulty equipment and devices can be detected and fixed.

Conclusions

This paper presents a case study where we successfully deployed an advanced control platform that consists of a cloud component and a physical component installed on-site, at 10 buildings at a university college in California. We were able to acquire real time data from the energy meters, BAS points and approximate occupancy count data through the count of Wi-Fi connected devices. Using this data, we were able to deploy FDD applications and ASO applications, producing energy savings up to 30% and 25% respectively. However, the COVID-19 pandemic and ensuing lockdown resulted in limiting cloud-based control, incomplete retro-commissioning efforts and impacted the measurement and verification of the performance of this control. Even though appropriate changes in the M&V methodology were developed to evaluate the performance of this platform, longer evaluation is required to obtain an accurate estimate of the savings.

References

ASHRAE. (2018, February 28). ASHRAE's BACnet Committee, Project Haystack and brick schema collaborating to provide unified data semantic modeling solution. Retrieved June 14, 2022, from https://bit.ly/2QwCtFf

Balaji B, Bhattacharya A, et. al. 2016. Brick: Towards a Unified Metadata Schema For Buildings. In *Proceedings of the 3rd ACM International Conference on Systems for*

Energy-Efficient Built Environments (*BuildSys '16*). Association for Computing Machinery, New York, NY, USA, 41–50. DOI:https://doi.org/10.1145/2993422.2993577

California Energy Commission. (2018). *2019 Building Energy Efficiency Standards*. California Energy Commission. Retrieved June 10, 2022, from https://www.energy.ca.gov/programs-and-topics/programs/building-energy-efficiency-standards/2019-building-energy-efficiency

Chen Z, Jiang C, Xie L. Building occupancy estimation and detection: A review, Energy and Buildings, Volume 169, 2018, Pages 260-270, ISSN 0378-7788, https://doi.org/10.1016/j.enbuild.2018.03.084.

Clark, C.; Prakash, A.; Pritoni, M.; Kloss, M.; Gupta, P.; Nordman, B., et al. (2020). Harvesting the low-hanging fruit of high energy savings -- Virtual Occupancy using Wi-Fi Data. *Lawrence Berkeley National Laboratory*. Retrieved from https://escholarship.org/uc/item/45s9h8v0

Haystack. *Home – Project Haystack*, 2022. https://project-haystack.org/.

Kramer H., Lin G., Curtin C., Crowe E., Granderson J. 2020. Proving the Business Case for Building Analytics. Lawrence Berkeley National Laboratory (2020), 10.20357/B7G022

Kluyver, T., Ragan-Kelley, B., et.al. (2016). Jupyter Notebooks – a publishing format for reproducible computational workflows. In F. Loizides & B. Schmidt (Eds.), Positioning and Power in Academic Publishing: Players, Agents and Agendas (pp. 87–90).

MelRok Energy IoT. *Get started - melrok energy IOT - learn how to use our tools today*. Select your gateway: Melrok Touch. 14 June, 2019. Retrieved March 25, 2022, from https://melrok.com/get-started/#selectgateway

Mills, E, and Mathew, P. 2009. "Monitoring Based Commissioning: Benchmarking Analysis of 24 UC/CSU/IOU Projects". United States. https://doi.org/10.2172/961599.

O'Grady, T, Chong, H, Morrison, G. A systematic review and meta-analysis of building automation systems, Building and Environment, Volume 195, 2021, 107770, ISSN 0360-1323, https://doi.org/10.1016/j.buildenv.2021.107770.

Pritoni, M.; Paine, D.; Fierro, G.; Mosiman, C.; Poplawski, M.; Saha, A.; Bender, J.; Granderson, J. (2021). Metadata Schemas and Ontologies for Building Energy Applications: A Critical Review and Use Case Analysis. Energies, 14, 2024. https://doi.org/10.3390/en14072024

Pritoni M., Lin G., Chen Y., Vitti R., Weyandt C., Granderson J. (2022). From fault-detection to automated fault correction: an empirical study. Building and Environment, Volume 214, 2022, https://doi.org/10.1016/j.buildenv.2022.108900

Acknowledgements

This work was supported by the California Energy Commission under the Electric Program Investment Charge (EPIC) Program, Solicitation Grant Funding Opportunity Number: EPC-16-034, entitled “Automated Cloud-based Continuously Optimizing Building Energy Management System (ACCO-BEMS). LBNL is supported by the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. Thanks to Anne Eisele and Pomona College for hosting this project. We would also like to thank the project lead, Zero Net Energy (ZNE) Alliance (ZNEA); as well as our partners and collaborators -- MelRok LLC.