

Existing Wireless Infrastructure Can Be a Low-Cost Path to Occupancy-Based Commercial Building Control

Cecil Scheib, New York University

Sam Scudere-Weiss, Splunk

Douglas Reagan, Splunk

ABSTRACT

Occupancy-based building controls are a commonly used energy management technique in institutional buildings. With better insight into occupancy patterns, facilities and energy managers can lower equipment to part load and change equipment scheduling to match actual (instead of predicted, or generous) needs. Building schedules can be adjusted to allow for more "off hours", especially with the rise of remote and hybrid work scenarios. Task-specific hardware is available that allows for such control but adds first cost expense and creates another system for facilities or energy management departments to maintain.

The ubiquity of smart devices and centrally managed wireless access points in government, institutional, or large corporate settings, allow re-use of an existing network to determine occupancy with a reasonable degree of granularity as to location and time. By examining wireless access point connection logs and using tuned heuristics to distinguish "false positives" of multiple devices, always-on stationary devices, and IOT devices, an anonymous count of occupants in a space can be estimated and updated, without investment in new systems or infrastructure and in buildings that do not have swipe card or other entrance / egress occupant counting.

New York University (NYU) was able to "tame" the large-bandwidth access point connection data source using middleware to create a dashboard that could be used for meaningful schedule setting, and potentially, real-time HVAC control. NYU found that with some allowance for adapting wireless connection data for an unanticipated use, this can be a cost-effective way to approach occupancy-based energy management for emissions reductions.

Introduction

With over 53,000 students, New York University is the largest private university in the United States. It is also one of the largest employers in New York City, with more than 19,000 faculty, staff, and administrators occupying about 15,000,000 owned and leased square feet in 200 buildings spread across Manhattan and Brooklyn. NYU's first climate commitments included a 2007 pledge to reduce greenhouse gas emissions intensity from buildings 30% from the 2006 baseline in ten years, by 2017, and carbon neutrality by 2040. After the 30% reduction was achieved, NYU updated its goal to a net 50% reduction from the 2006 baseline by 2025. Buildings represent over 99% of NYU's Scope 1 & 2 emissions, as well as over 2/3 of New York City's (NYC 2022), so deep reductions in building energy use are required to confront the climate challenge. Due to its setting in a highly urban, fully built-out area, new construction and/or substantial building renovations are rare, meaning the bulk of energy reductions in existing buildings must come from operational savings.

Some of the most cost-effective of those savings come from careful consideration of building scheduling, ensuring that lights, HVAC, and equipment are turned down or off to the greatest extent possible in unoccupied spaces. In an academic environment, spaces have the potential to be used at almost any time, as faculty and students may perform study or research late at night and on weekends, in addition to during the business hours that might be expected in a commercial real estate scenario. Therefore, savings from schedule adjustments are difficult to realize unless facilities managers and space users agree on schedules of occupancy. Bringing data to the discussion of space usage can be very informative and help conversations reach a mutually agreed upon conclusion. Referring to observed data on patterns of occupancy allows those seeking energy reduction opportunities to counter claims that “our space is used 24-7” and directs efforts to where the greatest energy savings can be obtained with the least disruption to current space use patterns.

The term “Great Energy Disconnect” has been used to describe the fact that while office buildings showed starkly lower occupancy during Covid, similar levels of energy reduction have not been seen. Part of this is inherent in the ability of building systems to run at part load, and also due to legal arrangements between owners and tenants. However, lack of control of building systems based on actual occupancy is also part of the problem. Greater visibility into actual occupancy patterns is essential to achieving cost-effective and durable greenhouse gas reductions as the US commercial market potentially moves into an ongoing remote/hybrid work situation (UGC, 2020, and Margolis and Rauch, 2021).

Background and Context

A survey of current active projects around reducing consumption for operational infrastructure by focusing on adjusting schedules of centralized heating and cooling across buildings or campuses is not new. Based on our understanding of the available literature, most of those other initiatives require either major alterations to existing HVAC systems and/or the addition of sensor equipment (Asim et al. 2022). Some other installations have leveraged Wi-Fi infrastructure to determine occupancy levels to adjust schedules as well, but most required additional investments in custom software and modifications to those systems that might not be acceptable to all schools or organizations (Clark et al. 2022). Other systems have used integration with lighting controls (Myer and Sandahl 2020, Hackel et al. 2020), which NYU has also done to a limited extent. However, lighting occupancy controls are less likely to be completely standard across a large university or corporate campus, at least one with buildings of various uses and vintages.

Business and Cost Drivers

After making its first climate commitment in 2007, NYU reduced energy usage about 5% year-on-year, with the largest reductions likely from a campus-wide reconsidering of building schedules and hours of operation. Many of NYU’s buildings are connected to one of multiple building management systems (BMS) that can be controlled from a central operations center. The rapid speed and scale of reductions encouraged NYU to achieve its 30% reduction goal, promised by 2017, in 2012 – five years early. However, a walk-through of campus buildings in the middle of the night, on weekends, during holidays, or “edge” hours just before and after typical occupied periods will still reveal well-lit, temperature controlled, but unoccupied spaces.

NYU has estimated potential emissions reductions from building schedule adjustments to be 5-10% of the needed total reductions to meet its 2025 climate goal.

Since BMS schedule adjustments are low cost, the energy savings from schedule resets show good paybacks, little or no need for additional funding, and an overall strong business case. Funds freed from powering unoccupied buildings can be redirected towards the University's academic mission. On the other hand, as an academic institution, there is often partial occupancy throughout the night, on weekends, and on holidays in lab and research spaces, faculty offices, and student study and gathering areas, and this needs to be considered. In addition, schedules in some types of academic buildings – for instance, those that house the performing arts – may vary greatly based upon performance schedules and rehearsals.

Project Considerations

Mix of Automation and Human Intervention

There are many options for standalone systems that measure building occupancy and adjust HVAC accordingly. NYU uses several systems that can provide this type of functionality, from simple room-based occupancy lighting sensors to more complex demand control ventilation and “body counters” in larger meeting rooms, classrooms, and theaters. The inexpensive, switch- or ceiling-mounted lighting sensors used by NYU do not transmit or store occupancy data, and developing a campus-wide network of sensors that do, or deploying a new campus-wide system across 15,000,000 square feet of built area, would be costly and likely not generate sufficient savings to justify the expense.

Leveraging Existing Infrastructure

There already exists an interconnected network that extends to essentially every space owned, leased, or managed by NYU: its wireless network for connecting laptop and desktop computers, phones and tablets, and other “Internet of Things” devices. Leveraging this existing network could allow for occupancy data gathering without deployment of a parallel or additional network. To allow this, a system for processing the large amounts of connectivity data from the network had to be developed, including acquiring, storing, parsing, and displaying occupancy data.

Occupant Satisfaction and Privacy

To ensure adoption and continuation of the building scheduling program, the NYU community of students, faculty, staff, and administrators had to be satisfied with their ability to access spaces that were appropriately scheduled. This meant extensive conversations with about two dozen schools, institutes, and business units. Each school or unit's occupancy needs were discussed as a narrative, putting academic needs first while simultaneously emphasizing the University's affordability, cost savings, and climate goals.

Privacy was a key concern and was given primary consideration from the inception of the program through deployment and regular use. From the beginning, the project requirements were that no personally identifiable data of any kind be stored or transmitted using the occupancy dashboard. The only data available would be the number of people on a certain floor of each building at a given time. There would be no way to determine from the dashboard who those

people were, any other information about them, or even their role within the university (for instance, student or employee). All that would be seen was a count of unique, connected devices determined to be representative of human occupancy by the middleware platform.

Supporting Technology

Current Campus Infrastructure

Reliable, ubiquitous, and continual connectivity is a nonnegotiable need in today's academic environment. NYU has made a substantial investment in its current wireless network, including over 10,000 Cisco Wireless Access Points covering almost 200 buildings across its New York City campus in Manhattan and Brooklyn. Those access points connect back to Cisco Wireless Lan Controllers (WLC) that centralize the management of user authentication and device connectivity. In addition, Cisco's Cloud Based Solution, DNA Spaces was also installed and used as the primary component to provide access to the data needed to support this initiative. Cisco's DNA Spaces offloads the processing load associated with data access from the controllers. Secure data access to other applications is provided via an Application Program Interface (API). This enabled the project to leverage existing available technology on campus without introducing any additional load on the infrastructure.

Data Capture, Analytics and Visibility Middleware

DNA Spaces provides large amounts of data on connections to the Wireless Access Points; however, significant data processing is required to onboard, process, analyze and provide visibility to that data from the Cisco infrastructure. NYU had been working with Splunk, a company that provides an enterprise data analytics platform, for other cyber security and systems management functions across much of NYU's current information technology infrastructure. A feasibility discussion determined that Splunk's platform would be flexible enough to analyze and present that data usefully.

Analytics and Data Visibility

Data Flow

As a middleware solution, Splunk was already taking in wireless data from the Wireless LAN Controllers. However, the project team decided that it would be better to pull data from Cisco's DNA Spaces Cloud based solution via an API, since the API allowed the Splunk Dev Team to query only the data fields required to support this use case which limited the already voluminous data coming into the WLAN Controllers. Those data fields included the time, building, and a unique identifier. Data would then flow into a dedicated Splunk Index (Figure 1), which is a free-flowing data store where the data would be kept in raw log form.




Figure 1: Flow graphic.

Speed and Privacy

After data hits the Splunk index, the Splunk program would run a process against that data that would create a summary of the data called a summary index. The summary would be used as the primary source of data for the dashboards used by the team. That summary index process sped up the search functions from the dashboards significantly, reducing load at the time of the search. Initial implementations of the dashboard could take up to an hour to load and display usable data, whereas later iterations show data in seconds.

In addition, while transferring that data, there was a sub-process run to transform any personal identifiable fields from the records to protect user privacy. But there was still a need to have a unique identifier that Splunk could use to reference that association of the device with the Access Point. To do that, during the sub process, splunk ran a hashing algorithm against the User ID. That non-reversible hash function (SHA512) cannot be inverted to show the original NetID and therefore there is no risk of revealing anyone's identity (Figure 2).




Figure 2: Privacy via SHA512 hash.

Accurate User/Occupant Counts

In a typical campus environment, there are many IP connected devices that are not carried by a person and associated with an occupant who might come and go from a space. This might include projectors, desk phones, desktop computers and other devices. The team needed a way to determine if a device was being carried by a user and filter out those other static devices from the summary data index.

This is done via an algorithm that uses movement between APs over a segment in time to determine static devices and filter those devices out of the counts. The main challenge from this approach was making sure devices “jumping” between two access points because of their location were not counted as active users. To solve these challenges the algorithm checks that a device has connected to at least three access points within a 24-hour period and adds the device to a list of active devices. The anonymized, hashed, unique identifier is used to ensure that a user with multiple devices is only counted once. This list is then utilized in the dashboards referenced below. The one drawback of this approach is that a user without an authenticated device will not be counted by the solution. However, the data seems to be sufficient for monitoring building usage trends.

Dashboards – Persona Based Visibility and Realtime Analysis

A project requirement role-based dashboards that provide the ability to analyze and visualize occupancy levels over a variety of time frames, by building or by area (school, department, business unit, etc.). The dashboards were designed to be used by energy or sustainability staff without needing a background in information technology software or database programming. Reports were needed to document occupancy levels while working with the University stakeholders to jointly decide on changes to BMS schedules.

The campus dashboard gives a top-level view of all buildings on campus connected to the occupancy system. Time ranges can be manually set (with a broad range of presets available) to view overall occupancy by campus (since NYU has several clustered locations in Manhattan and Brooklyn) and building. It’s a starting place to analyze key metrics, including average and minimum daily occupancy and number of hours under pre-determined occupancy level thresholds. Figure 3 shows an example of the all-campus dashboard.




Figure 3: All-campus dashboard.

To drill down to an individual building, users can access a building dashboard that summarizes occupancy on a floor-by-floor basis. Time ranges can again be set manually or using presets. This dashboard (Figure 4) is best used to analyze the number of occupants by floor over time and can be used to determine low occupancy levels at certain times of day and week on each floor. Cisco’s system, based on DNA Spaces data, uses proprietary algorithms to determine device location and provide an assessment of occupants by floor. This seems to be adequate to reduce confusion from incorrect floor counts (perhaps a device connected to an access point one floor above or below) or even incorrect building counts (where adjoining buildings share a common wall), perhaps because both access points and modern devices prioritize the strongest signal and NYU’s commercial buildings, with concrete floors and brick walls, reduce transmission between adjacent spaces. Theoretically, it is possible to use these algorithms to break down the data at a level even more granular than floor by floor, but that would not seem to add useful functionality to the system while adding complexity. It also presents potential concerns about privacy.




Figure 4: Individual building dashboard.

In addition to an engineering focus on individual buildings, setting schedules benefits from a view to how buildings are managed, which at NYU is often on the school or unit level. A school or a large administrative unit within the university at large may occupy part or all of several buildings, with a single point of contact for energy management and sustainability staff in setting schedules. Therefore, we developed a school/unit dashboard that grouped buildings or floors based on areas with common administrative control (Figure 5).




Figure 5. Occupant counts based on school or unit, by floor, for the NYU College of Dentistry (COD).

To assist in working with stakeholders, a trends dashboard shows granular level of occupancy across campus, highlighting times of day (or specific days) where occupancy levels fall below certain pre-determined thresholds. These thresholds are set with respect to absolute occupancy (trying not to heat/cool a building with just a handful of occupants) as well as relative occupancy (five people on a floor that has a full capacity of twelve is different than five people on a floor that has a full capacity of 100).

The trends dashboard (Figure 5) can show granular trends of occupancy level across two buildings. It can be useful to provide justification for a school or administrative unit as to the

reasoning and benefits behind proposed schedule changes.




Figure 6. Trends dashboard allows pre-processing of actionable data.

These dashboards will continue to evolve over time as any gaps from the current process emerge.

Data Measurement and Validation

After dashboard development, staff walked some floors in some buildings and did manual counts of occupants based on visual inspection to the best degree possible. This was then compared to dashboard values to validate the accuracy of the algorithm and data presentation. While the numbers never agreed exactly, they matched to within acceptable limits, especially since the data is not being used for live or real-time control. Pomona College has also documented a useful protocol for “ground truth” counts (Clark et al. 2020).

Using the Results

At this time, the dashboards are used retrospectively, helping to determine opportunities for building schedule changes by selecting an appropriate time of the academic year and time of day or week for analysis, as shown in Table 1.

Table 1: Potential building schedule strategies, additional to existing night and holiday schedules

Strategy	Time of academic year	Time of day or week	Theory of change	Potential HVAC hours of savings
Weekday nights	Mid-semester	9pm-1am	Classes may be held until 9 or 10pm, with a few occupants remaining in the building, given the impression that the building is occupied until midnight or later. However, past the end of classes, very few occupants may be present, and may not require continued presence in the same building.	1-2 hours per night, 4 or 5 nights per week, 35-40 weeks per year 140-400 hours/year total (up to 4.6% of year)
Weekday mornings	Mid-semester	6am-9am	It may be believed that occupants arrive as early as 7 or 8am, requiring HVAC starts of 5am or 6am. However, occupancy may be extremely light before a later time, perhaps 9am.	1-2 hours per morning, 4 or 5 nights per week, 50 weeks per year 200-500 hours/year total (up to 5.7% of year)
Weekends	Mid-semester	9am-9pm	NYU buildings are at perhaps 10% of peak occupancy on Sundays, somewhat higher on Saturdays. Buildings with very limited occupancy could have HVAC settings adjusted.	12 hours per day, 1 or 2 days per week, 40-50 weeks per year 480-1200 hours/year total (up to 13.7% of year)
January term (J-term)	Late December - late January	9am-9pm	Some buildings may be essentially unused during J-term, allowing for much greater setbacks.	12 hours per day, 7 days per week, 4-5 weeks per year 336-420 hours/year total (up to 4.8% of year)
Summer term	12-14 weeks, late May - end of August	9am-9pm	Some buildings are much more heavily used in summer than others. Some functions, such as summer classes, may be able to be consolidated into certain buildings, allowing greater setbacks elsewhere.	12 hours per day, 7 days per week, 12-14 weeks per year 1008-1076 hours/year total (up to 12.3% of year)

Note: savings in right-hand column are not cumulative as some time periods considered overlap

Using this data, predicted savings from potential actions can be estimated. In Figure 7, this is shown as annual greenhouse gas emissions reductions in MTCO₂e (“Carbon Savings (Tons)”), the annual economic value of energy savings alone, and the annual economic value of all savings, including reduced costs for security personnel in closed buildings. The scenarios shown are similar to but not exactly the same as those shown in Table 1.




Figure 7. Cost and emissions reductions from various schedule reduction scenarios.

Once general occupancy levels are taken from the dashboard, a heuristic – part routinized, part based on experience – is used to determine opportunity. For instance, the absolute occupancy numbers are only meaningful in relationship to floor size and number of occupants. If the dashboard shows eight occupants at midnight, that could be half of a small team on a 2,500 square foot floor of a narrow building, or it could be a basically empty 35,000 square floor in a building with a large floorplate.

The slope of the curve is also valuable – occupancy drops rapidly in office spaces and in classrooms, where groups of people tend to have consistent schedules and arrive and leave about the same time, and less obviously in research and other student and faculty spaces, where schedules may vary greatly. Perhaps there is a way to automate some of this heuristic in the future.

Once opportunities have been identified, occupants or leadership for a given space can be approached for a discussion about changing building schedules. Sometimes answers might be obvious: we assumed a library was used all night, but data consistently shows zero occupants between 2am and 6am. In other instances, there might be more of a discussion: should we keep a 100,000 square foot building open past midnight for a handful of occupants? What are their needs, and is there some other place on campus already planned to be scheduled occupied during that time where their needs could be met?

Challenges and Lessons Learned

Initially, NYU attempted to perform the analysis of data from the Cisco wireless networks with internal resources. Conceptually, this was straightforward. However, the large amount of data, combined with the complexity of the algorithms used to determine actual occupants, as opposed to stationary connected devices or connections to the network from outside the building, made this infeasible. A middleware solution with greater analytic capability was critical towards overcoming this, or more accurately, to allowing NYU personnel to manage the project instead of implementing it.

Ensuring consistent building and floor naming conventions took a significant amount of time and effort. Probably common to many urban campuses, buildings at NYU may have multiple addresses. A building that occupies a city block may have four or more legal addresses, at least one on each side. The electric, gas, water, and district steam utilities may all use different addresses, and NYU has 1800 utility accounts of this type. These may or may not correspond to the address the building is commonly known by. Further complicating matters, the building may have a common name, and that common name may not have a consistent spelling in building lists (for instance “3rd St”, “3rd Street”, “Third St”, and “Third Street” may all be found). For historical reasons, NYU does not have a single “source of truth” for unique building IDs. As a result, humans can parse NYU building lists, but automated lookups will generally fail unless custom lookup tables or “crosswalks” correspond one name domain to another. With almost 200 buildings in the analyzed set, this was non-trivial.

Analysis by academic or business unit was another naming challenge. NYU does maintain a list of which unit occupies or manages a space, but the coding is not necessarily consistent. For instance, for historical or system limitations, a space might be coded “PROVOST” or “PVST”. Again, custom lookup tables were required to generate meaningful dashboard data that could be used to approach building occupants and discuss schedules in a way that corresponded to their understanding of what spaces were under their control.

For others looking to replicate this work, ensuring the naming conventions of the buildings in DNA Spaces is correlated with the common names being used in other areas may be a critical task.

While the team was able to avoid the costs and complexity of a separate hardware network for occupant estimation by using the wireless network, the flip side of doing so is that a non-standard use of the Cisco DNA Spaces data may rely upon data persistence or retention qualities that were not anticipated in the original system design and specification. An 8-hour gap in data in one building, perhaps due to a transmission gap between different parts of the Cisco system, may be acceptable for wireless connectivity logging. As a result, there is not necessarily built-in redundancy or automated data validation for that need. However, for occupancy data purposes, an 8-hour gap in data can be material in determining building schedules, and Murphy’s law seemed to dictate that the data gaps occurred right where we most wanted to look. Since data was not being cached indefinitely, gaps in the transfer from Cisco DNA Spaces to the Splunk backend not caught within a few days were lost. The team spent significant effort tracking down and mitigating the causes of these data gaps, even though they had not previously been material to reliable operation and uptime of the NYU wireless network.

Maintaining an emphasis on privacy from the beginning of the project is critical. In at least one case, campus community concerns about privacy protection have led to a public apology on behalf of the institution, and the destroying of collected data (Wrighton 2022).

Conclusion and Next Steps

NYU's wireless-network-based occupancy data portal has been used successfully to adjust building schedules and reduce hours of occupied operation in the BMS, saving energy, cost, and emissions in a cost-effective way. For example, a schedule reduction pilot during January term (beginning in December 2021) in nine buildings reduced emissions by 37 tons MTCO_{2e} over just 12 days. Scaling across the NYU campus (200 buildings, though not all are subject to schedule changes, such as laboratories and residences) and to different low-occupancy time periods including summer session and holiday weeks, the savings could be far larger. NYU is expanding this program across campus during summer 2022.

In the future, some of the analysis might be more automated, saving staff time in running queries and considering the results. There might be more of a feedback loop, showing users occupancy against the schedules they requested and automatically highlighting opportunities for further changes. Probably deeper savings could be obtained with this strategy (Balaji et al. 2013), but whether the increased benefits justify the increased complexity and risk for error is still to be determined.

While the connection of this system to the BMS was beyond the scope of this project, the potential for future real-time control remains. Various issues, from cybersecurity to data validation, uptime, and algorithm reliability, as well as hardware/software engineering for the linkage between quite disparate systems, would have to be addressed before this occurs.

Acknowledgements

The authors gratefully acknowledge the work of Syed Raza and Ellie Tam (NYU IT) and Jauna Vitale (NYU Office of Sustainability) for their efforts to successfully implement this system.

References

- Asim, N., Badiei, M., Mohammad, M., Razali, H., Rajabi, A., Chin Haw, L., & Jameelah Ghazali, M. (2022). *Sustainability of Heating, Ventilation and Air-Conditioning (HVAC) Systems in Buildings-An Overview*. International journal of environmental research and public health, 19(2), 1016. <https://doi.org/10.3390/ijerph19021016>.
- Balaji, B., J. Xu, A. Nwokafor, R. Gupta, and Y. Agarwal. 2013. Sentinel: Occupancy Based HVAC Actuation using Existing WiFi Infrastructure within Commercial Buildings. Roma, Italy. SenSys.
- Clark, C., A. Prakash, M. Pritoni, M. Kloss, P. Gupta, B. Nordman, M. A. Piette, M. Kamel, T. Semann, A. Eisele, D. Hage, P. Flannery. 2020. "Harvesting the low-hanging fruit of high energy savings -- Virtual Occupancy using Wi-Fi Data." In *Proceedings of the 2020 ACEEE Summer Study on Energy Efficiency in Buildings* 12:121–135. Washington, DC: ACEEE.
- Hackel, S., Cowan, C., Zhou, J., Wolgamott, C., and Myer, M. 2020 "Lighting controls: not just for lighting anymore!" .” In *Proceedings of the 2020 ACEEE Summer Study on Energy Efficiency in Buildings* 3:103-120. Washington, DC: ACEEE.

Margolis, S., and Rauch, M. 2021. *The Great Energy Disconnect*. Albany, NY: NYSERDA (New York State Energy Research and Development Authority).
<https://www.nyserda.ny.gov/All-Programs/Energy-Efficiency-for-Commercial-Tenants/Resources/The-Great-Energy-Disconnect>.

Myer, M., and Sandahl, L. 2020. “Integrated and Advanced Lighting Controls: The Next Frontier in Energy Savings?” In *Proceedings of the 2020 ACEEE Summer Study on Energy Efficiency in Buildings* 3:212-227. Washington, DC: ACEEE.

Nikdel, L., Janoyan, K., Bird, S., and Powers, S. 2018. *Multiple perspectives of the value of occupancy-based HVAC control systems*. Building and Environment (IF6.456), DOI: 10.1016/j.buildenv.2017.11.039.

NYC (New York City). 2022. *GHG Inventory*. Accessed March 2022.
www1.nyc.gov/site/sustainability/reports-and-data/ghg-inventory.page

UGC (Urban Green Council), 2020. *The Great Energy Disconnect*. New York, NY: Urban Green Council. <https://www.urbangreencouncil.org/content/events/great-energy-disconnect>.

Wrighton, M. 2022. *Message Regarding Data Analytics Pilot Project*. Accessed March 2022.
president.gwu.edu/message-regarding-data-analytics-pilot-project.