# Using Wi-Fi Location-Based Services (LBS) for Commercial Building Occupancy Sensing

*Matt Kirk, Design.Garden*
*Di Sui, P.E., Ramboll Group*
*Robert Lysholm and Lester S. Shen, Ph.D., Center for Energy and Environment*
*Ben Bartling, CEM, CMVP, Lee Shaver, and Scott Hackel, Slipstream*
*Matthew Schwartz, Parallel Technologies*
*Brad Kult, P.E. (MN), HGA*

## ABSTRACT

From May 2019 through October 2022, this DOE-funded project investigated and demonstrated the use of Wi-Fi Location-Based Services (LBS) to perform occupancy sensing in commercial buildings. Wi-Fi LBS can be used to detect the presence of Wi-Fi enabled mobile devices and laptops that accompany occupants as they move through the building. These signals can be used to determine occupant presence, head count, and location. When integrated with the building automation system, this emerging technology approach can be used to manage other connected systems such as lighting and HVAC to reduce energy usage in the building and improve occupant comfort. An open source location detection algorithm was developed, which uses data collected from three or more Wi-Fi access points to determine the presence and estimate the location of mobile devices and laptops. Access points can detect Wi-Fi enabled devices even if they are not connected to the existing Wi-Fi network. Building occupancy is determined based on the presence, location, and movement of these devices through the space. From lab and small-scale in-situ testing, the Location Detection Algorithm (LDA) was found to be accurate to within 10 feet and could be further refined by tuning the algorithm for the specific space characteristics such as layout and obstructions (walls, furniture, etc.). An open source method to integrate the occupancy data with existing building automations systems was investigated. The Wi-Fi occupancy sensing approach was then demonstrated and validated at two commercial buildings located in Minnesota and Wisconsin.

## Introduction

Location-based services (LBS) have been used successfully in a variety of applications in healthcare, retail, and hospitality buildings to track and in some cases interact with individuals. These services include wayfinding, asset tracking, marketing, information, and push notifications. Specific applications include the following.

- *Wayfinding* in shopping malls (Mall of America, 2017) and in the office (Castellanos, 2020)
- *Asset tracking* applications in manufacturing (Reddy, 2015), retail (Swedberg, 2019), and healthcare (Yoo, et al., 2018)
- Proximity *marketing* in stores where advertising content such as coupons or offers is provided wirelessly through apps to shoppers when they approach specific items and brands in the aisle (Mittal, 2019)
- *Information* provided in self-guided museum tours (Chun, 2016 and Pau, 2017)
- Location-based *push notifications* (MacFarlane, 2019)

LBS use real-time locating systems (RTLS) that identify and track the location of objects and people through wireless technology. Within the defined space, fixed reference points (sometimes called beacons, readers, or access points depending on the technology used) receive signals transmitted by the devices or tags that the object or person is carrying. The tag's signal strength to a specific reference point is a measure of proximity and, using multiple beacons/readers/access points, the location of the tag can be determined through trilateration.

Wi-Fi-enabled devices can serve as a tag to provide location since they constantly seek Wi-Fi networks that are in range. Wi-Fi access points (APs) that are in range receive the media access control (MAC) address and received signal strength indicator (RSSI) of each device whether that device is logged on to the Wi-Fi network (i.e., associated with the network) or not. The MAC address is a hardware identification number that uniquely identifies that specific device on a network. For privacy, some mobile devices transmit randomized MAC addresses over time to anonymize the device. The RSSI is a measure of the power level of the signal that is being received by the AP. Once a device is within range of an AP, device presence and identity are immediately sensed and location can be determined from the RSSI via trilateration. Occupant activity might also be inferred by the space where a Wi-Fi-enabled mobile device is located and/or by the movement of that mobile device around the space.

There are other systems available on the market for location detection, such as IR cameras, RFID tags, Bluetooth beacons, and people-counting sensors. However, they either require installation of additional servers and/or devices or require occupants to carry special devices. Wi-Fi RTLS are well suited for occupant sensing in commercial buildings because:

1. A majority of the occupants will be carrying a mobile device on them or within their close proximity. For instance, in higher education a factor of 2.5–3 devices per person could be used to account for a student's cell phone, laptop, and tablet/other devices. These device count estimates are commonly used by network engineers when planning for wireless network capacity in higher education classroom buildings.[1]
2. Tracking does not require the individual to carry additional hardware (like a badge for RFID beacons) or to have a native app loaded on their device and connected to the network (like Bluetooth beacons).
3. Existing Wi-Fi access points can, by default, engage all Wi-Fi enabled smartphone users (iPhone or Android), whether they are logged onto the Wi-Fi network or not.

Minimal additional infrastructure is required since existing Wi-Fi hotspots (installed for the communications needs of cell phone and computer users) provide the coverage that can be used for the location tracking. As long as the Wi-Fi system can provide the MAC address and RSSI data to perform the trilateration calculations, Wi-Fi LBS can be performed in the space.

Wi-Fi RTLS provides the following capabilities that make it a good candidate for occupant sensing:

1. There is virtually no latency between the time when a new device has entered the room and when an updated occupancy count is calculated (as opposed to $CO_2$ sensors). The frequency of the count is based on the time interval chosen to rescan the AP data.
2. Failure rates that occur with motion sensing are avoided when the device is at rest.

---

[1] B. Kult, personal communication, 2019.

3. Wi-Fi LBS can map mobile devices across rooms and other areas of interest and trigger customized operating conditions. If individuals choose to opt in, Wi-Fi LBS can identify them and further customize their experience.

Wi-Fi RTLS is limited by the resolution of the occupant's sensed position. The research generally agrees that a resolution of ±10 feet is achievable indoors.[2] This resolution is acceptable when occupants reside well within a single area of interest, but presents a challenge when occupants reside near the edge of two or more areas, as the Wi-Fi RTLS system cannot confidently resolve the appropriate area of the occupant. Nonetheless, this level of resolution is well matched to the area (500 to 3,000 square feet) typically served by a single zone in a building HVAC system.

There are commercially available Wi-Fi platforms that provide Wi-Fi location analytics. The two main manufacturers are Cisco with their Meraki Location Analytics[3] and Connected Mobile Experiences (Cisco CMX) Analytics[4] and HP with their Aruba Analytics and Location Engine (ALE).[5] The Cisco and HP platforms provide their location data through application programming interfaces (APIs). This data can then be made available for use by building automation systems (BASs) or Internet of Things (IoT) platforms that reside inside the building to take action based on the estimated building occupancy and the predicted locations of mobile devices detected by the wireless platform.

The company Sensible Building Science (SBS)[6] employs Wi-Fi RTLS occupancy detection using the Cisco CMX wireless platform. This system includes applications of the latest RTLS methods to provide energy efficient building operation. The SBS Bridge software receives real-time Wi-Fi data using the location analytics provided by the network platform, specifically the Cisco Connected Mobile Experiences (CMX) software solution. Originally tested in buildings on the University of British Columbia campus,[7] SBS is piloting their environmental control systems in buildings in Canada, the United Kingdom, and the United States. As of 2017, SBS has used their approach in over one million square feet of commercial and institutional space, serving over 100,000 occupants in real-time. Preliminary results indicate an average annual savings of 5% in whole-building energy use and have found that buildings with variable occupancy and demand control ventilation (DCV) offer the greatest potential for savings. In lecture halls with periodic classes, SBS was able to reduce fan runtime by 20%–40%.

This DOE-funded project seeks to investigate and field validate the use of access point-based Wi-Fi Location-Based Services (LBS) in commercial buildings. This paper will report on the following tasks of the project.

- Developing a non-proprietary, open source algorithm to perform trilateration to predict the location of Wi-Fi enabled devices.
- Testing and validating the LDA to provide occupancy data.

---

[2] B. Kult, personal communication, 2019.

[3] https://meraki.cisco.com/solutions/location-analytics

[4] https://www.cisco.com/c/en/us/td/docs/wireless/mse/10-6/cmx_config/b_cg_cmx106/the_cisco_cmx_analytics_service.html

[5] https://www.arubanetworks.com/products/location-services/analytics/ale/

[6] https://sensiblebuildingscience.com/

[7] https://news.ubc.ca/2017/03/30/innovative-software-converts-wi-fi-data-into-energy-savings/

- Planning the field demonstrations of the LDA in commercial building settings.

## Developing the Wi-Fi LBS Location Detection Algorithm (LDA)

The Location Detection Algorithm (LDA) calculates the distance between a Wi-Fi enabled device and Wi-Fi access points (APs) based on the strengths of three or more received signals. Within the defined space, Wi-Fi-enabled devices that the person is carrying transmit signals searching for Wi-Fi networks to join. The measure of the signal power level is the Received Signal Strength Indicator (RSSI) in dBm. The RSSI of the Wi-Fi device to a specific AP reference point is a measure of proximity and, using multiple APs, the location of the Wi-Fi device can be determined through trilateration. In Figure 1 below, the three reference points denoted by $R_1$, $R_2$, and $R_3$ each read a signal ($s_1$, $s_2$, and $s_3$, respectively) from the Wi-Fi device. The location of the Wi-Fi device or person P is where the radial signal strengths $s_1$, $s_2$, and $s_3$ intersect.

*Figure 1. Determining Location via Trilateration*

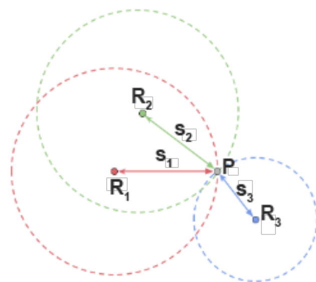


The distance from each AP is calculated from the RSSI based on the Hato-Okumara model.

### Hato-Okumara Model

The Hato-Okumara model is a power signal-based position method defined by the equation (Bose and Foh, 2007):

$$log\ d\ =\ \frac{1}{10n}\left(P_{TX} - P_{RX} + G_{TX} + G_{RX} - X_a + 20\ \ log\ \lambda\ - 20\ \ log\ \ (4\lambda)\right) \qquad (1)$$

where:

$d$ is the estimated distance between the transmitter (Wi-Fi router or access point) and the receiver (Wi-Fi-enabled device).
$P_{TX}$ (dBm) is the transmitted power level. For the OpenMesh APs used in this study, this is 13 dBm.
$P_{RX}$ (dBm) is the power level measured at the receiver or the measured RSSI.
$G_{TX}$ (dBi) is the antenna gain of the transmitter. For the OpenMesh APs, this is 2.5 dBi.
$G_{RX}$ (dBi) is the antenna gain of the receiver. For the OpenMesh APs, this is 2.5 dBi.
λ (m) denotes wavelength of the signal and can be estimated to be 0.12m for the middle frequency of the 802.11b channel (2442 MHz).
$X_a$ is a normal random variable with a standard deviation of a and is in the range of 3 dB to 20 dB, depending on the building construction and any obstacles that will reflect, diffract, or scatter the signal. (standard deviation = ±5 dBm)

$n$ is a measure of the influence of obstacles like partitions, walls, and doors. For an unobstructed line-of-sight (LOS) path between the transmitter and receiver, $n$ = 2 but for obstructed paths, n should be between 4 and 5.

Grouping all the defined constants into one term A, equation (1) becomes:

$$log\ d = \frac{1}{10n}(A - P_{RX}) \tag{2}$$

where A can be defined as the reference signal strength received in dBm:

$$A = P_{TX} + G_{TX} + G_{RX} - X_a + 20\ \ log\ \lambda\ - 20\ \ log\ \ (4\lambda) \tag{3}$$

The distance d is therefore determined by:

$$d = 10^{\frac{(A - P_{RX})}{10n}} \tag{4}$$

Bose and Foh (2007) suggest using a multi-model approach to improve accuracy. From their empirical data, they split the signal propagation model into two parts based on proximity to the access point. At closer ranges (< 5m) they suggest a higher value for the n factor with n = 5, and for distances greater than 5m, n should be set to 4. In the lab environment, this distance corresponds to -49 dBm. So, for RSSI > -49 dBm, n = 5 and for RSSI < -49 dBm, n = 4. They also suggest that a multi-model approach could be used for differences between a LOS and non-LOS environment.

**Location Detection Algorithm (LDA)**

An application was coded in JavaScript to analyze the measured signal data (RSSI and MAC address) collected from APs to (1) identify Wi-Fi enabled laptops and hand-held devices in a space, (2) determine the location of each device, (3) correlate the device(s) to an occupant, (4) assign the occupant's location to a zone in the space, and (5) Estimate the headcount for the zone. The estimated headcount from the algorithm could be used by building automation systems (BASs) to operate the building's lighting and mechanical heating, ventilation, and air-conditioning (HVAC) systems. The LDA was developed using OpenMesh Wi-Fi APs that allowed us to retrieve the MAC addresses and RSSI values they read.

The MAC address of the Wi-Fi-enabled device is a unique device identifier primarily assigned by the device's manufacturer. Often, the MAC address will contain a vendor portion or the OUI (Organization User Identifier) which will be the first six digits of the MAC address. These can be compared to a library of published vendor MAC addresses to exclude Wi-Fi devices that will not be of value to occupancy sensing, such as manufacturers of Wi-Fi enabled office equipment.

To provide identity protection and prevent tracking, many mobile devices can perform MAC address randomization. The device's operating system will replace its MAC address with randomly generated values. The randomization can be performed periodically over time. Regardless, if the LDA collects a nonsensical MAC address due to randomization, this device can be inferred to be a mobile device and tracked for occupancy sensing.

The LDA performs trilateration using the RSSI data measured from three or more APs and visually displays the calculations' results within the floor plan of the space where the APs reside. Figure 2 shows a sample visualization from simulated data in a space with four APs locating one Wi-Fi device.

*Figure 2. Test visualization of trilateration using four APs*

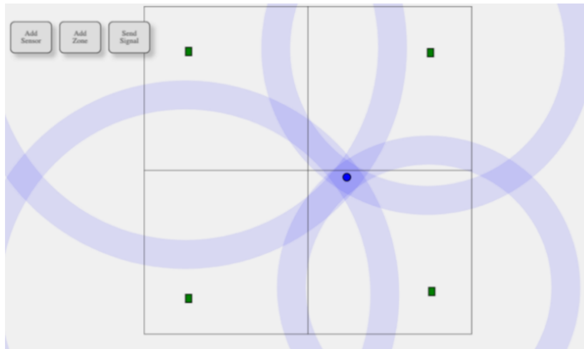

The green squares show the location of the APs and the blue rings around each AP show the calculated radial distance the target device (blue dot) is from the respective AP, the distances calculated from the measured RSSIs. The thickness of the ring shows the probability distribution of the signal's origin, which in this case is the expected measurement error calculated as ± the standard deviation of RSSI readings over the specified time span of the measurements (one hour in our tests). With four APs to perform the trilateration, the overlap of the four rings gives the likely locations, and where all the rings intersect is the most likely location of the target. The value of n was the same for each AP and the space was simulated with unobstructed LOSs for each AP.

The APs receive signals from Wi-Fi devices not only within the space of interest but also outside the space in adjacent rooms or outside the building. Two rules were created to identify the devices to track within the defined space. The algorithm will only track signals that (1) are received by all the APs in the space and (2) pass a defined signal strength threshold. Testing found that location accuracy increased with the number of signals received from the Wi-Fi device, with a minimum of five readings to provide good accuracy. With these filters, the algorithm was tested in the field to determine its ability to accurately calculate occupant count within a zone.

## Algorithm Testing and Validation

Three sites were used to perform in-situ testing and validation. Because of the pandemic, these sites had limited occupancy, allowing us the flexibility to place known and identifiable devices in the spaces to simulate occupants and target individuals who were present in the spaces. Occupant presence, counts, and location were verified by self-reporting, visual observation, and digital video footage from the temporary deployment of Blink home security cameras.[8] The three sites were the Design.Garden offices, the CEE Lending Center, and the Parallel Technologies' Innovation Lab.

### Design.Garden Offices

Seven OpenMesh OM2P APs were installed in the Design.Garden offices and three Wi-Fi devices (three Google Homes) were placed at various locations in the space. Figure 3 shows the predicted Wi-Fi device locations mapped onto a blueprint of the space, visualized by the LDA. The green squares show the locations of the APs, the small circles show the locations of

---

[8] https://blinkforhome.com

each Wi-Fi device, and the larger circles show the locations of the devices predicted by the algorithm. The lines connecting each set of circles show the distance errors of the calculations.

*Figure 3. Locations of APs and Wi-Fi Devices in Design.Garden Office Space*

Over five days, 256,286 probe requests were collected. These probe requests were not evenly distributed across devices and time; nonetheless, each AP collected at a minimum one probe request per 15 minutes from each device.

Using this data, we first attempted to recreate the experiment outlined in Bose and Foh (2007), using their RSSI-to-Distance transformation adjusted for our equipment and with an arbitrarily chosen $X_a$ of 5 dBM. The transformation operates on individual RSSIs; however, for our experiment we calculated an average RSSI using a simple moving average (SMA) to mitigate noise. Using a sampling width of 4 hours in our SMA, we produced transformations every 15 minutes over 5 days, giving us an array of distance guesses per sensor and device through time. Comparing our distance estimates to actual values showed an average error of 7.7 feet. Bose and Foh (2007) report an average error of 7.5 feet, hence we feel we have successfully recreated their experiment.

Next, we used our trilateration algorithm to transform the distance guesses into two-dimensional location guesses per device through time. Comparing our location guesses to actual values produced an average error of 10 feet. These initial results confirm the viability of Wi-Fi LBS, and they nearly meet our stated goal of a ten feet maximum error.

We had based our stated goal of ten feet on results we had generally seen in literature regarding the one-dimensional RSSI-to-distance transformation; however, we discovered through experimentation that the error in two-dimensional trilateration varied significantly with placement of devices and sensors. This effect is inherent to trilateration and is profound enough that we could conceivably arrange the devices and sensors so that the error would never exceed 10 feet, even if given incorrect RSSI readings.

Therefore, meeting our goal of ten feet is not enough to prove that we have successfully implemented a Wi-Fi LBS. Instead, to prove success we should show that we are significantly outperforming a failing Wi-Fi LBS given the same placement of devices and sensors. We define a failing Wi-Fi LBS to be one that fails to correlate RSSI readings to location, and we define the resulting error to be the failure threshold of a system given the same device and sensor placement.

Finally, we define the confidence of a Wi-Fi LBS system to be the difference between one and the ratio of its location guess to its failure threshold, expressed as a percent. Confidence near 0% represents an inability to perform Wi-Fi LBS, while far less than 0% represents a calibration error.

To calculate our failure threshold, we replaced the RSSI readings for Day 1 with randomly generated values, then we reran our trilateration experiment. This system meets our definition of a failing Wi-Fi LBS; therefore, we can calculate our failure threshold as the difference between its location guesses versus actual values. Comparing our average error for Day 1 to our failure threshold gives our Wi-Fi LBS a confidence of 49%, with an improvement of 9.1 feet over the error of a failing Wi-Fi LBS.

Having quantified our confidence in our Wi-Fi LBS, we looked for opportunities to improve our accuracy by optimizing our use of the RSSI-to-distance transformation. Notably, the transformation includes a variable, $X_a$, to account for random path-based attenuation. We assumed $X_a$ would be normally distributed around a non-zero mean per device and sensor pair, and that we could estimate this mean value over time given the distance. This estimated mean could then be substituted for $X_a$ in all future transformations to yield more accurate location guesses.

Unfortunately, this method assumes the path between the sensor and device does not change over time, so it seems unsuitable for Wi-Fi LBS. However, we assumed that if we kept the sensor position static while moving the device throughout all reasonable locations, the resulting estimated mean, $X_\mu$, could then be substituted for $X_a$ regardless of device location.

Therefore, we developed an algorithm to estimate $X_a$ per sensor using three of our devices as reference points representing a subset of all reasonable locations, then we reran our trilateration experiment for the fourth device, Device 2. Compared to our previous location guess for Device 2, we improved our error from 6.7 feet to 3.4 feet and improved our confidence from -13% to 82%.

This analysis suggests that the use of permanent fixed reference points will be useful not only to commission the RTLS but also to maintain the accuracy of the Wi-Fi RTLS by providing a means to continuously calibrate the system. Inexpensive Wi-Fi devices can be placed within a space to serve as Wi-Fi beacons. A number of wireless, connected devices that typically exist in offices can also serve as stationary reference points, such as: Wi-Fi-enabled television displays in conference rooms, office equipment such as printers and copying machines, cameras, sensors, networked lights, and the APs themselves.

Because the LDA allows us to select which APs can be used in performing the trilateration calculations, we investigated how the placement of the APs affected the accuracy of the device location prediction. By choosing the number and locations of the APs, we could investigate how the position of the Wi-Fi devices relative to the APs affected the accuracy of the location predictions. This suggests that the ideal placement for APs to provide good location predictions will be along the perimeter of the space being monitored. This is different from the typical AP placement for good Wi-Fi coverage in a space that would place the APs more centrally. More central positioning of the APs also minimizes the number of APs needed to provide the wireless network while a minimum of three APs is needed for location prediction. Placing the APs along the perimeter to provide both good Wi-Fi coverage and accurate location predictions will likely require more APs to be installed than if only Wi-Fi coverage is required.

**CEE Lending Center**

The Lending Center is occupied by CEE staff who approve home improvement and energy efficiency loans for qualified Twin Cities residents. Loan closings also take place in the Lending Center. Four OpenMesh APs were placed along the perimeter of the Lending Center. Data from the APs were uploaded to the Design.Garden database server for analysis by the LDA. Wireless battery powered Blink home security cameras were placed within the Lending Center to create a visual record of occupancy in the space and to verify the calculations of the LDA. These cameras are Wi-Fi enabled and used as reference points. Figure 4 shows the floor plan with the locations of the APs (denoted by the gray rectangles).

*Figure 4. Floor Plan of the CEE Lending Center and Reception Area.*

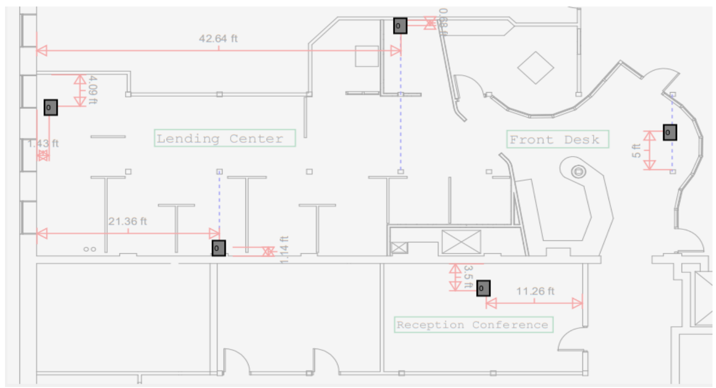


Testing of the LDA in the CEE Lending Center allowed us to determine the accuracy of the LDA to track motion and occupancy count compared to the information collected from the Blink cameras. The photographic evidence also allowed us to determine the accuracy of occupant count based on the number and location of mobile devices (phones, tablets, and laptops) detected.

Even though the pandemic greatly reduced occupancy at the CEE Lending Center and the administrative offices, the Lending Center will typically have one or two staff working in the office per day and there is one administrative staff person who will be present at the front desk. Applicants for energy and home improvement loans also come to the office for closings, which take place in the small meeting room in the Lending Center. CEE staff who do occasionally come into the office will enter through the doors at the front desk and then pass to the administrative offices. So, some traffic will be detected passing through this area.

First, we estimated Xμ for the experiment using two reference devices and two stationary target devices. We substituted Xμ for Xa in the LDA, i.e., we tuned the LDA, then we continued to collect data to validate our Xμ. Over six days, our tuned system maintained an error less than 10 feet when locating the two stationary target devices.

Having validated our system, we then performed a controlled experiment to test the LDA's ability to track motion. In the experiment, one known device was moved fourteen times throughout five rooms in the CEE Lending Center over four hours. In our previous lab experiment, we used an SMA with a sampling width of four hours to mitigate noise; however, in this experiment the minimum duration spent resting before each move was five minutes, hence we were limited to a sampling width of five minutes.

The LDA located the moving device with an average error of 14.2 ft. The average error decreases to 12.2 ft. when we first average the samples of each room separately, and to 9.7 ft. when we also exclude the room that is furthest from our reference devices; therefore, we conclude that (1) our estimated $X\mu$ is insufficient for the entire space. The average error decreases from 14.2 ft. to 12.6 ft. when we remove the worst LDA result from each room; therefore, we conclude that (2) the short sampling width is insufficient to mitigate noise.

The LDA operated with high confidence (>50%) for 90 minutes of the four hour experiment, and with any Confidence (>25%) for only half of the experiment. Our LDA ran with high confidence (and an average error rate of 5.9 ft.) for 78.3% of the experiment when excluding the room that is furthest from our reference devices. Therefore, we conclude that the LDA can confidently track the motion of devices in a space given reasonable exceptions.

Finally, we validated the Wi-Fi LBS's ability to detect presence by comparing our location guesses to 308 minutes of activity detected by security cameras over three days. The LBS detected 230 minutes (74.7%) of the activity and detected the correct count of people over 142 minutes (46.1%).

**Parallel Technologies' Innovation Lab**

The Parallel Technologies, Inc. (PTI) Innovation Lab is used for product demonstrations, testing, and training. Because of the pandemic, the office area was typically vacant and most staff worked remotely. The Innovation Lab did have some occupant traffic due to deliveries, technicians entering and leaving through the area, staff performing work in the lab, and some traffic to the adjacent warehouse.

The existing PTI Wi-Fi network had one Cisco Meraki AP placed in the lab area, two Meraki APs in the adjacent warehouse, and one Meraki AP in the adjacent office area. For the test, seven OpenMesh APs were placed at locations in the lab, warehouse, and office area. Figure 5 shows the floor plan of the Innovation Lab (the area denoted with the 5,780 sq. ft. floor area) and warehouse space (the area denoted with the 5,258 sq. ft. floor area) with the locations of the OpenMesh APs denoted by the green dots and the locations of the Meraki APs denoted by red circles (where three of our APs were also placed).

*Figure 5. Floor Plan of the Parallel Technologies Innovation*


Security cameras had been installed in the lab and CEE staff were able to use the recorded footage from these cameras to compare actual occupant counts and locations with those predicted by both the LDA and the Meraki location analytics at the corresponding dates and times. Two cases were performed for the validation.

For the first case, two Wi-Fi-enabled devices with known MAC addresses (a laptop and a tablet) were placed at specific locations in the space and the accuracy of the LDA predictions was evaluated. The LDA consistently predicted the devices' locations within our error margin of 10 ft.

Having validated our system, next we validated the Wi-Fi LBS's ability to detect presence by comparing our presence guesses to presence detected by security cameras 25 times over two days. The LBS detected presence 16 times (64%) and detected the correct count of people over 13 times (52%). Because of the sparse occupancy in the space due to COVID restrictions, an error margin of a couple of devices would result in no occupant presence detected and greater percent error.

For the second test, the LDA was used to monitor occupancy during an open house that took place in the Innovation Lab. On October 28, 2021, Parallel Technologies hosted a client appreciation event in the Innovation Lab from 2:00 p.m. to 6:00 p.m. Approximately 80 to 100 people were in attendance. CEE staff were in attendance to document the occupant count, location, and traffic. They also brought multiple Wi-Fi-enabled devices with known MAC addresses. This allowed us to compare the LDA predictions with the devices' known locations at various designated times during the event.

First, we tested the LDA's ability to provide headcounts of large groups of guests (with unknown MAC addresses). The PTI lab includes several static Wi-Fi devices, so we began by identifying and excluding MAC addresses seen in the days before the event, then we confirmed a zero device count both before and after the event. Next, we compared headcounts performed by

CEE staff throughout the event to device counts calculated by the LDA using sampling widths of 5, 15, 30, 120, and 240 minutes.

As expected, we found that (1) lower sampling widths produced lower counts, at worst counting only 11% of the guests, whereas (2) higher sampling widths failed to acknowledge waning attendance, at worst counting 195% of the guests. In our opinion, a sampling width greater than 30 minutes is too long for responsive HVAC control, while a sampling width of 5 minutes is too short to count all guests. Therefore, we feel that a sampling width of 15 or 30 minutes is best, on average counting 30% and 42% of guests, respectively.

*Figure 6. Headcounts versus device counts (by sampling width) in the PTI Innovation Lab.*

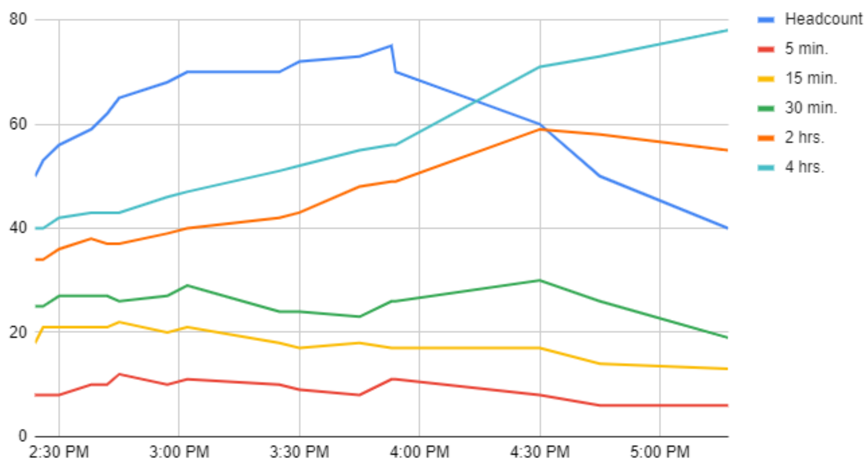


## Field Demonstrations

Two commercial office sites were recruited to evaluate the performance of the LDA for occupancy sensing. At the time of this writing, we have begun monitoring these sites and are in discussions with two additional demonstration sites. We have installed OpenMesh APs for Wi-Fi device detection. We have also installed cellular routers onsite to upload the data from our APs directly to Design.Garden's server for processing with the LDA. Creating the cellular hotspot for our APs allayed any cybersecurity concerns that the respective IT department had about the project accessing the site's wireless platform.

### Slipstream Office Building, Madison, Wisconsin

The Slipstream Wisconsin offices are located in a two-story office building in Madison, Wisconsin. The testing is being performed in their second-floor offices where five Ubiquiti APs support their Wi-Fi network. We placed eight of our OpenMesh APs in the space (denoted by green dots), five at the same locations as the Ubiquiti APs (red circles) and another three to provide additional coverage for evaluating the LDA. The floorplan of these offices is shown in Figure 6.

*Figure 6. Floor Plan of the Slipstream Offices Showing the Locations of the APs.*

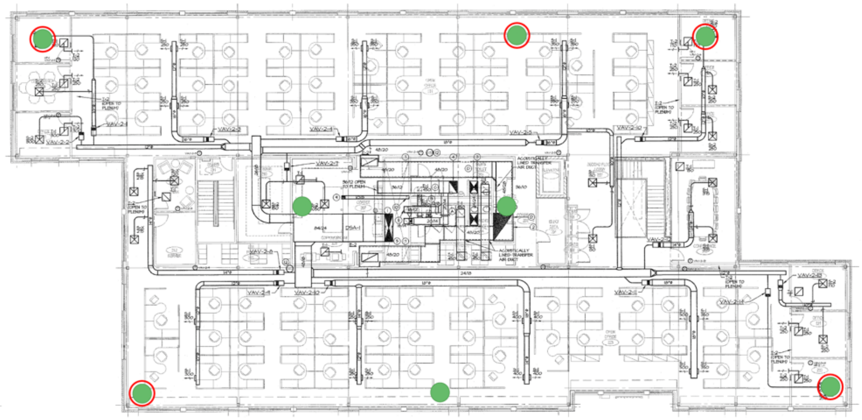

**Minnesota Administration Building, Saint Paul, Minnesota**

The Minnesota demonstration site is in the offices of the Minnesota Department of Administration Materials Management and Plant Management Divisions located in the Administration Building at the State Capitol Complex. For the demonstration, we placed eight OpenMesh APs at the same locations as the existing APs to replicate the existing Wi-Fi coverage. Figure 7 shows the locations of the APs on the floor plan of the offices. The green dots represent our OpenMesh APs and the red circles represent the locations of the existing APs.

*Figure 7. Floor Plan of the Minnesota Department of Administration Offices Showing the Locations of the APs.*

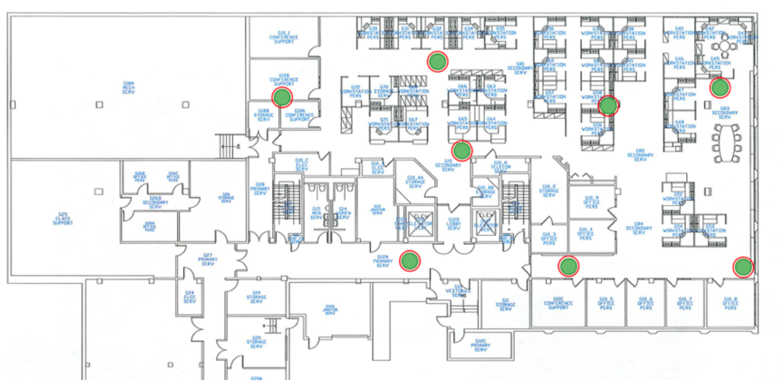

**Occupancy Validation**

A number of approaches will be used to validate the occupancy of the two sites. The principal method used at both sites will be to enlist an on-site person to log the occupant headcount with dates and times in specific rooms or spaces at the site. This data will then be used to validate the corresponding data calculated by the LDA. We were not allowed to install the Blink cameras at the Admin Building for video verification.

At the Slipstream site, we will also use video collected from Blink cameras to monitor occupancy. Rather than having staff examine the video to collect occupancy data, computer vision will be employed to analyze the video of occupants moving in and out of the entrance of a room or space.

Slipstream will also analyze the key card security system used to gain entry to the second-floor offices to validate occupancy load. The key card data that is collected accounts for the door openings as people access the space. The key card data will be an approximate count of the occupants that enter the second-floor offices but does not account for those that leave or if multiple people enter when the door is opened by a single pass of the key card.

Furthermore, Slipstream will perform occupancy sensing tests using a computer vision app that is installed on a laptop tethered to a USB camera. Computer vision will attempt to classify people in the video frame and then use object tracking algorithms to count the people crossing through the video feed from the USB camera. A GitHub repository[9] has been created to represent the computer vision code in an open source format. The app running the computer vision will also have a REST API endpoint for an IoT device (VOLTTRON at the Slipstream office) to log the data.

**BAS Integration**

The VOLTTRON IoT platform has been implemented at the Slipstream office and integrates the Trane HVAC controls via BACnet/IP protocol. VOLTTRON is an open-source distributed control and sensing software platform. It is an ideal platform for us to integrate our Wi-Fi LBS occupancy data and test its use with building operation. The occupant data calculated by the LDA will be provided to Slipstream's VOLTTRON edge device via an API call over the Internet on 60-second intervals to the Design.Garden server. VOLTTRON will use this information to instruct the BAS to condition the offices' defined zones based on the level of their respective occupancies via BACnet write commands created by VOLTTRON "control agents." The VOLTTRON platform supports data acquisition (i.e., BACnet reads of the 60-second interval data by the BAS) via VOLTTRON "drivers" on a variety of protocols, as well as the creation of Python scripting to control HVAC via BACnet write command. In the VOLTTRON nomenclature these are called VOLTTRON "agents."

Two zones will be tested: an open office area with three VAV boxes and a conference room with a dedicated VAV box. The control sequences of the VAV terminal boxes that were originally set up by the HVAC control contractor will be overridden with a new control sequence via BACnet by VOLTTRON based on head count data. Slipstream will log data in an SQL database onsite and display the data on a Grafana dashboard to monitor the performance over time. Based on headcount data retrieved from Design.Garden, the new HVAC sequence will be as follows.

(1) For both the open office area and the conference room, the setpoint temperatures will be changed to the setback conditions when there are no occupants in the zone.
(2) The conference room VAV box damper will be closed when there are no occupants in the zone. This damper control is only to be applied to the conference room since it is in the building's core area.
(3) When the building is unoccupied, the VOLTTRON platform will release all BACnet overrides implemented on the HVAC controls to allow for original unoccupied space temperature control sequences to operate.
(4) The space will run in original occupied mode when occupants are sensed in the zone.

---

[9] https://github.com/bbartling/building-people-counter

## Discussion

During the installation of the access points at the two demonstration sites, we learned that Datto, who purchased OpenMesh, would no longer support the CloudTrax platform that allows us to download our data from our OpenMesh AP devices. We found that the Open Mesh APs that we used with CloudTrax continue to provide data to us but we could not register any new devices to the platform. Consequently, we collected a set of CloudTrax-compatible APs for the Slipstream site and will use this platform as long as it is available to us. For the remaining OpenMesh APs that we had, we found that PlasmaCloud[10] offered a cloud management tool that would provide the functionality that CloudTrax previously provided with the OpenMesh APs. For the Minnesota Administration site, we converted a set of OpenMesh APs to the Plasma Cloud platform. Our database server is able to collect data from both platforms and the LDA has no issues using data from either source or from other platforms such as the Cisco Meraki location analytics API.

We are currently monitoring the two demonstration sites and are working to finalize two additional sites. We plan to present the results of these demonstrations during our oral presentation at the 2022 ACEEE Summer Study. The end date of this project is September 30, 2022 and the final report will be published during the Fourth Quarter of 2022.

## Wi-Fi LBS as an Energy Efficiency Strategy in Buildings

The current work has shown that the Wi-Fi LBS LDA is capable of detecting occupant presence and head count through the Wi-Fi-enabled devices they are carrying with them. This information can then provide an estimate of the headcount of occupants in a space. If permission is granted, it can also identify the occupant. Table 1 defines the types of occupancy sensing information that Wi-Fi LBS enable along with the most likely areas of application for use as inputs that can save energy in commercial buildings.

*Table 1. Occupancy Information Resolution Levels. (Wang, et al., 2019)*

| Resolution Level | Functional Definition | Technical Definition | Application |
|---|---|---|---|
| Occupant presence | Is the space occupied? | Are the number of qualified Wi-Fi devices greater than 0? | Lighting, HVAC schedule optimization |
| Occupant count | How many people are in the space? | How many qualified Wi-Fi devices are in the space? | HVAC control optimization: Demand control ventilation, model predictive control; Energy benchmarking, M&V, Fault detection diagnostics |
| Occupant identity | Who is the person? | Who owns the device? | Personalized work environment management |

[10] https://www.plasma-cloud.com/

## Concluding Remarks

This paper describes our efforts to evaluate the use of Wi-Fi LBS for occupancy sensing by developing an open source algorithm that allows us to locate and count Wi-Fi enabled devices. We determined the location accuracy of the algorithm (within 10 feet) and identified how to tune the algorithm to achieve even greater accuracy with the use of fixed reference points. The work we will be performing at the demonstration sites will let us (1) evaluate the ability of the LDA to estimate occupancy within a space, (2) develop an approach to implement the data with a BAS, and (3) define the control sequences that a BAS would employ using the data.

## References


Bonthu, B., and M. Subaji. "An effective algorithm to overcome the practical hindrance for Wi-Fi based indoor positioning system." *Open Comput. Sci.* 2020. 10:117-123. https://doi.org/10.1515/comp-2020-0010.

Bose A., and C.H. Foh. "A practical path loss model for indoor Wi-Fi positioning enhancement, in 2007 6th International Conference on Information." *Communications and Signal Processing*. ICICS. 2007. 10.1109/ICICS.2007.4449717.

Chun, R. (2016, May 5). *The SFMOMA's New App Will Forever Change How You Enjoy Museums*. Retrieved December 9, 2019, from http://www.wired.com.

MacFarlane, K. (2019, April 8). *18 Inspiring Location-based Push Notification Examples & Ideas*. Retrieved January 16, 2019, from http://www.taplytics.com.

Mittal, S. (2019, June 19). *Proximity Marketing Examples: 28 Retail Companies Nailing it with their Campaigns*. Retrieved December 9, 2019, from http://beaconstac.com.

Pau, S. (2017, April 19–-22). *Audio That Moves You: Experiments With Location-Aware Storytelling In The SFMOMA App.* Retrieved December 9, 2019, from http://mw17.mwconf.org/.

Reddy, A. (2015, April 19). *The Growing Use of RTLS by Manufacturers.* Retrieved January 7, 2020, from http://www.rfidjournal.com.

Shen, L.S., D. Sui, C. Plum, and M. Kirk. (2020). *Wi-Fi Locations-Based Services for Occupancy Sensing in Buildings: A Technical Overview*. Center for Energy and Environment. Retrieved May 31, 2022, from https://www.mncee.org/sites/default/files/report-files/Wi-Fi-LBS-Tech-Overview.pdf.

Swedberg, C. (2019, December 30,). *Concept Store Delivers Product Content via NFC, RFID*. Retrieved January 7, 2020, from http://www.rfidjournal.com.

Z. Wang, T. Hong, M.A. Piette, and M. Pritoni. "Inferring occupant counts from Wi-Fi data in buildings through machine learning", *Building and Environment*, 158 (2019) 281–294.

Yoo, S., S. Kim, E. Kim, E. Jung, K.H. Lee, and H. Hwang (2018). *Real-time location system-based asset tracking in the healthcare field: lessons learned from a feasibility study.* BMC Medical Informatics and Decision Making. 18, 80.