Data-Driven Monitoring and Optimization of Classroom Usage in a Smart Campus

Thanchanok Sutjarittham
University of New South Wales
Sydney, Australia
t.sutjarittham@student.unsw.edu.au

Hassan Habibi Gharakheili
University of New South Wales
Sydney, Australia
h.habibi@unsw.edu.au

Salil S. Kanhere
University of New South Wales
Sydney, Australia
salil@cse.unsw.edu.au

Vijay Sivaraman
University of New South Wales
Sydney, Australia
vijay@unsw.edu.au

ABSTRACT

Student enrollments world-wide are increasing each year, while lecture attendance continues to fall, due to diverse demands on student time and easy access to online content. The resulting underutilization of classrooms entails cost penalties, especially in campuses where real-estate is at a premium. This paper outlines our efforts to instrument a University campus with sensors to measure classroom attendance, in a cost-effective and scalable manner without endangering student privacy. We begin by undertaking a lab evaluation of several approaches to measuring class occupancy, and compare them in terms of cost, accuracy, and ease of deployment and operation. We then instrument 9 lecture halls of varying capacity across campus, collect and clean live data on occupancy spanning about 250 courses over 12 weeks during session, and draw insights into attendance patterns, including identification of canceled lectures and class tests; our occupancy data is released openly to the public. Lastly, we show how classroom allocation can be optimized based on attendance rather than enrollments, resulting in potential savings of 52% in room costs.

CCS CONCEPTS

• **Information systems** → **Data analytics**; • **Computer systems organization** → **Sensor networks**; • **Hardware** → **Sensor applications and deployments**;

KEYWORDS

Classroom occupancy, Data analytics, Optimization

ACM Reference Format:

Thanchanok Sutjarittham, Hassan Habibi Gharakheili, Salil S. Kanhere, and Vijay Sivaraman. 2018. Data-Driven Monitoring and Optimization of Classroom Usage in a Smart Campus. In *Proceedings of ACM IPSN*. ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

ACM IPSN, April 11, 2018, Porto, Portugal

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM…$15.00
https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 INTRODUCTION

Higher education institutes continue to experience steady growth in enrollment demand [8]. A major factor limiting Universities in fulfilling this demand is real-estate, since enrollment in a course is capped by the capacity of the classroom to which the course is allocated. However, with recent trends towards student lifestyles that mix study with work and other commitments, as well as greater access to online content, there is ample anecdotal evidence that classroom attendance is often well below the enrollment number. This presents an opportunity for education institutes to better optimize the usage of classroom space based on attendance rather than enrollments. However, this requires real-time visibility into attendance, which can vary significantly between courses and across weeks of the session, so as to dynamically re-allocate courses to rooms while minimizing the risk of attendance exceeding room capacity.

Several methods are available to count the number of people in an indoor space, such as WiFi-based locationing, camera image processing, thermal imaging, ultrasound imaging, and beam counters affixed to entryways. Each method has its own pros and cons in the various aspects such as cost, power, communications, ease of deployment and operations, privacy, and accuracy. For example, WiFi-positioning and cameras endanger privacy; thermal and ultrasound imaging have low accuracy; and camera-based image processing is expensive. Furthermore, a method that works well in a small room may not be as effective in a large lecture theater, and cost/accuracy may also be impacted by the layout of the room, the number/width of doorways, and the availability of power and wired/wireless network connections. Understanding both the benefits and the challenges in measuring classroom attendance in a large campus of varying lecture-rooms cannot be done as a paper-study, and requires experimental field-work in deploying, integrating, operating and evaluating the various systems aspects in a real campus.

This paper describes our experiences in building a system for measuring classroom occupancy, and its deployment across 9 rooms in a large campus. We begin by testing several sensing methods in a lab environment, and characterizing their trade-offs in aspects such as cost, ease of installation, method of data extraction, and accuracy. We then make appropriate sensor selections, build a full system, and deploy it across 9 lecture-theaters of varying size across a university campus. We collect and clean the data to obtain real-time visibility of occupancy across these rooms in real-time over a period of 12 weeks, integrate it with University timetabling systems to infer attendance patterns of nearly 250 courses, and highlight interesting findings such as attendance trends, canceled lectures, and class tests. We also make our data openly available to the research community. Finally, we develop an optimization algorithm for allocating courses to rooms based on dynamic attendance rather than static enrollments, and show potential saving of 52% in room costs.

The rest of this paper is organized as follows: §2 describes relevant prior work. We present our lab evaluation of various sensing methods and their trade-offs in §3, while §4 describes our field deployment across campus and the interesting insights obtained therein. §5 presents our optimization technique and quantifies the benefits of attendance-based allocation, and the paper is concluded in §6.

2 RELATED WORK

An obvious approach to deducing occupancy counts is to use information from existing WiFi access point (AP) infrastructure to infer the location and number of users in a room [6, 12, 13, 18]. There are, however, factors that may affect the accuracy of people counting using this method – not every occupant may have a device connected to the AP, some may possess multiple devices, and devices carried by people outside the room may be connected to the same AP. Obtaining WiFi connectivity data may also constitute a violation of privacy if the identities of connected users can be deduced.

Several studies have also used video camera based approaches for people counting; [14, 15] have achieved good accuracy by applying complex image processing algorithms, however they require fairly heavy computational resources. The work in [19] successfully uses image processing to extract people counts from a large classroom with many occupants. Nevertheless, this method only works well when there is not much movement in the classroom. Privacy also remains an issue, especially if images and videos of people are taken without their explicit consent.

Some studies have used special-purpose sensors for occupancy counts – [9] deployed a complex sensor network including ambient-sensing and carbon dioxide monitors, and apply Hidden Markov Models to detect occupancy in offices. The method shows 73% accuracy, but is tested only in small rooms with less than 10 occupants. In [20] a single passive infrared (PIR) sensor is combined with intelligent machine learning to predict room occupancy; though the method offers a low-cost solution, it is tested in rooms with only 14 or less occupants. Work in [16] has collected the occupancy data a commercial space using depth sensors (Kinect for XBOX One) for a duration of 9 months, and authors have released their dataset in [17]. The work attempts to predict future occupancy using historical data, but it suffers from a high error rate of prediction (i.e. up to 2100%).

Much of the aforementioned works focus on dynamic occupancy detection as an end in itself, or for energy optimization (heating/cooling) purposes. The problem of allocating classrooms to courses in the University context has been studied by previous works such as [7, 10, 11]; however, they are based on (static) enrollment numbers, rather than (dynamic) attendance counts. To the best of our knowledge our work is the first to combine dynamic occupancy measurement with classroom space allocation. We however emphasize that the current work does not predict room occupancy (such models are the subject of future work), and instead shows the potential gains if the optimization were run off-line with *a priori* knowledge.

3 LAB TRIAL: SENSING METHODS AND TRADE-OFFS

We briefly describe the various sensing methods we tried in 3.1, and outline their relative trade-offs in 3.2, with a view towards making appropriate selections suitable for a larger-scale deployment across the campus.

3.1 Sensor Selection and Data Collection

Sensors: We investigated several commercial sensors and straightaway eliminated those that send data to the vendor's cloud servers, since we wanted to: (a) keep the data entirely on-premises and not risk it leaving our campus infrastructure; and (b) not be beholden to a vendor to access our own data, hence freeing us from ongoing service costs. In other words, we wanted a "sale" model of the device so we could have unfettered access to our data without any ongoing "service" fees. We were quite happy to buy spares of the units to cover for device failures; further, this model allows us to integrate data into a centralized repository to facilitate better analytics across the many data feeds we have on campus.

We narrowed our lab trials to four types of commercial sensors: EvolvePlus Wireless Beam Counter [2], EvolvePlus Overhead Camera [1], Steinel HPD Camera (pre-market release), and Steinel Presence Detector [3]. In addition, the University IT department provided us with timestamped connections logs from two WiFi access points (one inside our lab and one just outside), so we could compare our approaches to those obtained from WiFi logs. We note that the WiFi logs gave us personal user information such as their device MAC address, user-ID, and connection durations; we therefore obtained ethics clearance (UNSW Human Research Ethics Advisory Panel approval number HC17140) for this experiment.

The **Beam Counter** comprises a pair of infrared (IR) break-beam sensors mounted on the door frame, and counts the number of people passing through in each direction. It communicate the counts (for "in" and "out" directions) to a gateway every 30 seconds using a propriety wireless protocol, and the gateway then posts these readings via Ethernet to an SQL DB server hosted on a VM in our on-premises cloud infrastructure. The **Overhead Camera** is a thermal sensor mounted on the ceiling close to the entrance facing downwards, and counts the number of people passing below it. It also communicates the counts in each direction to the same gateway as the beam counter, which then forwards it on to the SQL-DB. We wrote a script that pulls data from the SQL-DB, stamps the data with the time and the unique UUID of the gateway, and posts as a JSON string to our master database (which holds data from many sources) via a REST API.

The **HPD Camera** (pre-market release) is a people counting sensor mounted in a corner with full view of the room. It uses in-built image processing to compute the number of people present within a configurable zone of interest. It is powered over Ethernet, and comes pre-configured with a server that be queried via a REST API. We wrote a "broker" script that polls the camera every 30 seconds to get the people count, and posts the time-stamped and sensor UUID-stamped data in JSON format to our master database. The **Presence Detector** is a passive infrared (PIR) sensor mounted on the ceiling in the middle of the room, and detects motion. Though it does not count the number of people in a room, it gives a binary indication on whether the room is occupied or not – this sensor can be used as a way to calibrate the other counting sensors which may accumulate errors with time. The PIR sensor sends its binary occupancy state every 60 seconds to its corresponding gateway via a propriety wireless interface, which then posts it to a broker script that again time- and sensor-UUID-stamps the data and posts to our master database.

Lastly, we receive a CSV file of daily WiFi connection logs for the two access points from our IT department every morning at 7am – real-time feed of data was not possible due to technical limitations of the AP vendor. We wrote a script to parse the log file and compute the number of unique users connected to each AP every 30 seconds – this was also posted to our master database.

3.2 Sensor Comparison

Our lab trial helped us compare the various counting methods in terms of their ease of installation, calibration, power and communications requirements, accuracy, cost, and privacy, as summarized in Table 1.

Our comparison across these measures is qualitative rather than quantitative. Even aspects such as accuracy, that can be quantified, depend on factors like room size and layout, mounting position, number of doors, and width of doorways, which can vary widely across deployment environments. We therefore resort to qualitative measures (low, medium, and high) in this table, derived from our experience across the rooms we instrumented, and we back these up with several data points presented later in the paper.

Installation: The thermal camera, HPD camera, and PIR sensor needed professional installation by certified tradesmen, since each needed special mounting brackets and extra wiring for mounting on (or near) the ceiling. We could install the beam counter sensor easily by ourselves using two-sided adhesive strips on the door frame at around waist-height.

Calibration and Positioning: Sensor positioning is another key factor in our comparison. The thermal camera needs to be positioned at a certain height range (i.e. 2.2m - 4.4m) recommended by the manufacturer and close to the entrance allowing the best coverage to count everyone that passes underneath. This requirement makes it hard or impossible to use the thermal camera in very large lecture halls with high ceilings. Beam counters require to be mounted at around waist-height (too low causes each leg to get counted separately, and too high causes the swinging arms to get counted!). Once an appropriate height is chosen for the beam counters, doors of all classrooms need to be outfitted in the same way. The HPD camera needs prior configuration for zone of interest that can vary across rooms depending on the room size and the place at which the camera is mounted. The PIR sensor is positioned

Table 1: Sensors comparison.

	Installation	Calibration	Power	Comms	Accuracy	Cost	Privacy
Beam counter	easy	easy	battery	wireless	high	medium	✓
Thermal camera	hard	medium/hard	AC power	wireless	low	high	✓
HPD camera	medium	hard	PoE	Ethernet	medium	medium	✓
PIR sensor	hard	easy	AC power	wireless	binary	medium	✓
WiFi	existing	existing	existing	existing	low	0	✗

at the center of the room (on the ceiling) to have a symmetrical coverage over an area that can also vary across rooms depending on their seating arrangement.

Power and Communications: Provisioning power was challenging for the thermal camera and PIR sensor, since the campus has pre-built and fixed wiring only in certain locations in each classroom. Therefore our Facilities Management was required to supply new exterior wiring for these three sensors. The beam counters are battery powered (with stated battery life in excess of a year), and the HPD camera required a special PoE switch that provides Ethernet for both power and communications. The corresponding gateways for the beam counter, thermal camera, and PIR sensor were hidden inside a closet with available power and Ethernet.

Accuracy: We performed several spot measurements in our lab to extract ground truth on occupancy. We found that the beam counter is the most accurate among the four techniques. We note that the beam counter has very good accuracy when the the door is narrow, like in our lab. However, for a wider doorway its accuracy is worse, since it does not always capture individuals walking in/out side-by-side (this became more evident in our field-trial, described in the next section). We found the accuracy of the thermal camera to be very sensitive to mounting position and distance from the entrance. Moreover, since the door of our lab opens inwards, it was not very conducive for the overhead thermal camera (mounting it on the outside of the room was not an option as it was a busy corridor). The HPD camera tended to have a non-zero absolute count error, which made its relative error high when the number of people in a room is small (e.g. less than 10) and low when the number of people is high (e.g. more than 40). We could not test its accuracy scaling to larger counts as our lab can only accommodate around 40 people. Lastly, the people count derived from the WiFi access points was wildly inaccurate, because our lab is adjacent to a busy corridor and study space that is busy with students during regular hours, and we could not distinguish who was inside versus outside the room.

Cost: The beam sensors and PIR sensors are priced in the range of a few hundreds of dollars, while the cameras are in excess of a thousand. The beam counter and thermal camera both need a gateway to send their readings to the back-end server, and each gateway is priced in at nearly a thousand dollars. Bear in mind that each gateway can connect up to 20 sensors (though our deployment described in the next section maps at most 4-5 sensors to a gateway in large lecture theaters). The beam counters therefore end up as a more cost-effective solution than the cameras for large-scale deployment across campus.

Privacy: Among the four sensing techniques, WiFi clearly endangers students privacy as their IDs are visible (due to PEAP authentication their devices perform to connect with the campus WiFi network). The HPD camera does on-board processing and does not store or transmit any images of people (though it is possible to log in to it to view the current image), and can hence be deemed to preserve privacy. The beam counter and the thermal camera are truly privacy-preserving, since they can only sense the number of people passing through the doorways without sensing any private attributes of the individuals.

Summary: The trade-offs discussed above are summarized in Table 1. WiFi is not an option as it compromises privacy and is inaccurate. The cameras are eliminated as being expensive, difficult to install/position, and poor in accuracy (though we are considering them for open spaces that do not have doorways). The PIR sensor has only binary output, and is used for re-calibration rather than counting. We therefore decided on a larger-scale deployment of the beam counter, based on its relatively lower cost, easy deployment, high accuracy, and good protection of privacy. Our deployment in classrooms is described next.

4 FIELD TRIAL: DATA PROCESSING AND VISUALIZATION

We worked with campus staff to identify appropriate classrooms for a field trial, and picked 9 rooms of varying sizes, as shown by top two rows in Table 2. Some of the doorways to the lecture-halls posed a challenge as they were very wide, increasing the likelihood that multiple students walking out side-by-side get counted as one. The data collected over the first few days was manually verified (volunteers were used to do head-counts) so as to obtain ground-truth and calibrate the errors. In what follows we describe our methods for data cleansing, linking with class-timetabling information, processing, and visualization using a web-UI.

4.1 Data Processing

We compared two methods for deducing the occupancy based on the number of entries and exits at each door.

Method 1 – Room Occupancy: Our first (naive) method for deriving occupancy is to set it to the cumulative number of entries minus the cumulative number of exits across all doorways of a classroom. However, errors arise when students walk in/out in groups; though we reset counts to zero at midnight each day, errors accumulating during the day can become significant.

Method 2 – Course Occupancy: To reduce the errors accumulating during the day, we enhance our method by computing course attendance independent of each other by linking our sensor data with course timetables databases obtained from our University. We assume that students may enter the room up to 10 minutes prior to start of the scheduled lecture time, and may leave up to 10 minutes after the scheduled lecture time. Attributing each entry and exit to a specific lecture therefore allows us to compute attendance per-course, and errors are not carried over from one lecture to the next even if they are adjacent in time to each other.

Accuracy of Counting: To evaluate the accuracy of our counting methods, we obtained ground-truth information by having volunteers physically count attendance during the lectures. We collected a total of 50 samples covering 31 lectures over 4 days. The ground-truth samples were collected from 8 out of 9 classrooms in which the sensors have been deployed. Table 2 shows the average error of the computed occupancy using the two methods described above, applied to the various rooms. As expected, course-based occupancy computation yields lower errors (average: 12.71%) compared to room-based occupancy computation (average error: 30.60%). We also noted that the room-based method gradually built-up errors over the course of a day, whereas the course-based method had a stable error irrespective of time-of-day (since the errors do not accumulate). However, it should be noted that the course-based method requires access to timetabling information, which may not be generalizable to other environments.

Table 2: Measured error from ground-truth of occupancy.

	BUS105	BUS115	CLB7	CLB8	MatA	MatB	MatC	MatD	PhTh
No. of doors	1	1	4	3	6	3	2	2	4
No. of seats	35	53	497	231	472	246	110	110	369
Room-based	21.7%	34.2%	89.5%	26.3%	25.5%	16.3%	14.6%	16.9%	NA
Course-based	13.0%	17.3%	4.6%	16.1%	8.0%	9.1%	24.4%	9.2%	NA

Occupancy data: Our weekly dataset, computed using the Method-2 above, is openly available for download [4]. Each row in a CSV file represents the real-time measurement from beam sensors comprising *time-stamp*, *week* of semester, room information including *room name*, *number of doorways*, and *number of seats*, course information including *course-id* (we have intentionally obfuscated the actual names of courses), *course start-time*, and *course end-time*, sensor measurements including *count-in*, *count-out*, and computed number of attendance (i.e. *occupancy*). Note that count-in and count-out are available for the entire day (even during times with no lectures scheduled), whereas occupancy is available only when a course is scheduled.

4.2 Data Visualization

Tool: We developed a web-application to provide an intuitive UI for real-time occupancy monitoring, using Shiny (an R web framework) and the interested reader can see the interface live at [5]. The user can choose the date and classroom to view occupancy level in real-time along with the lecture capacity and number of enrollments. The interface also allows the user to view the attendance pattern of a course (by choosing from the course dashboard tab), as well as the utilization rate (number of attendees divided by the total number seats available for each classroom) for different time-slots.

Insights: Our UI provides some interesting insights into attendance patterns. Fig. 1 show our UI output for a chosen room (CLB8) on a specific day during week 4 of term. It shows the number of attendees (solid green lines), enrollments (dotted orange lines), and the capacity (dashed blue lines) for six lectures that are scheduled between 9am to 9pm. In general, attendance is seen to vary widely across courses, and can be in the range of 10-90% of enrollments; interestingly, we can see here that the lecture scheduled for 1-2pm has enrollment of 211, but close to zero attendance; this indicates that the lecture was probably canceled, leading to a waste of allocated classroom space on that day. This shows the benefit of our tool in quantifying space utilization that is otherwise largely unknown to facility managers.

Our visualization tool also allows us to track the attendance pattern of various lectures over the semester, as shown in Fig. 2 for three selected courses (we have obfuscated course names). Course-059 in solid blue lines exemplifies a (typical) gradual decline in attendance over the weeks. Course-188 in dashed orange lines shows a sudden rise in attendance in week-6, which is indicative of a class-test (typically held in week 6 or 7 of the 13-week term) - we

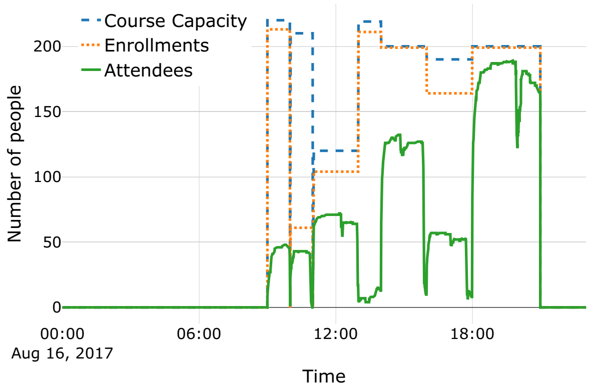

Figure 1: Occupancy pattern of a classroom.

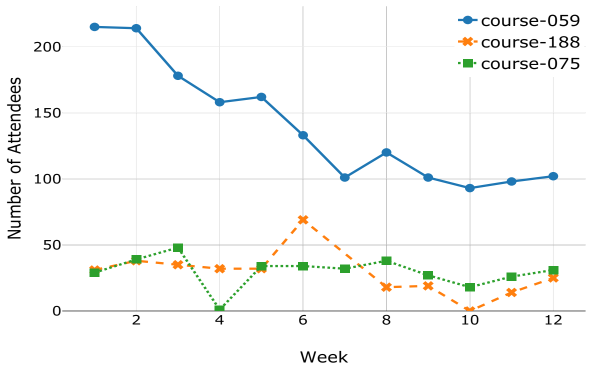

Figure 2: Attendance pattern of various course.

were able to verify this is indeed the case by looking up the mid-term test date on the course web-page. Lastly, Course-075 shown by the dotted-green lines depicts a canceled lecture in week 4.

Our tool also provides visualization of the utilization "heat-map" of the classrooms on a chosen day, as shown in Fig. 3 – bright (yellow) cells depict a utilization (ratio of attendance to capacity) approaching 1, while dark (blue) cells represent poor utilization closer to 0. For example, room MatB is under-utilized during the time-slot 3-4pm. Hovering over that cell, it shows that Course-031 has 15 attendees with the room capacity of 246 (i.e. 6% utilization). This interface helps campus managers track classrooms utilization, with a view towards more optimal allocation, as described next.

5 DYNAMIC ALLOCATION OF CLASSROOM

We observed in the previous section that course attendance varies over the weeks which leads to under-utilization of classrooms. This presents an opportunity for campus managers to employ a dynamic allocation scheme to save cost. In this section, we develop an optimization formulation to determine the potential cost savings. Our formulation assumes prior knowledge of attendance numbers, which is of course not practical, but is meant to establish a benchmark upper-bound on the potential cost savings (thus helping business stakeholders justify their upfront investment into instrumenting classrooms). A practical implementation could use historical attendance data (say from the previous year) to develop a dynamic schedule for a course using our optimization algorithm,

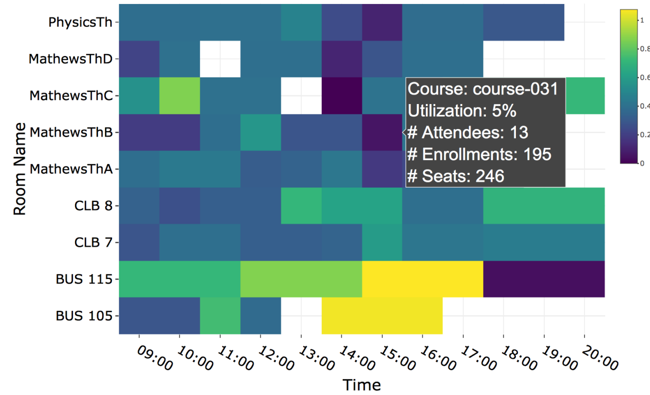

Figure 3: Occupancy Heatmap.

while leaving some margin for error arising from the use of retrospective rather than prospective attendance counts. The practicalities of this are beyond the scope of this paper and deserve a separate study in its own right.

5.1 Problem Formulation

We now formulate the problem of optimal classrooms allocation. First, let there be R rooms available on campus, and each room has a cost associated with it (proportional to its capacity). The cost of room j is denoted by C_j where $1 \leq j \leq R$.

We consider our optimization problem over a one-hour window. Therefore, let there be L courses in operation over the window of interest, and the number of attendees and enrollments for course i are denoted by e_i and a_i respectively, where $1 \leq i \leq L$.

We define our variable by $x_{i,j}$ that represents whether or not course i is allocated to room j over the window of interest. $x_{i,j}$ can be written in the following equation:

$$x_{i,j} = \begin{cases} 1 & \text{if } \textit{course } i \textit{ allocated to room } j \\ 0 & \textit{otherwise} \end{cases} \tag{1}$$

Therefore, the total cost of allocation for a given window is specified as:

$$J = \sum_{j=1}^{R} \left\{ C_j \sum_{i=1}^{L} x_{i,j} \right\} \tag{2}$$

Our aim is to minimize the total cost J in (2). Note that allocation of a course to a room incurs a full cost of that room, and an unallocated room incurs no cost.

We note that each course can only be allocated to one room during a window. These L constraints are captured by:

$$\sum_{j=1}^{R} x_{i,j} = 1 \ \forall \ i \tag{3}$$

Further, a room cannot be occupied by more than one course at a time. These R constraints are captured by:

$$\sum_{i=1}^{L} x_{i,j} \leqslant 1 \ \forall \ j \tag{4}$$

We need to ensure that attendees of a course are fitted into the allocated room. To allocate a course, we consider two cases: (a) enrollment-based, where the room capacity needs to be larger than

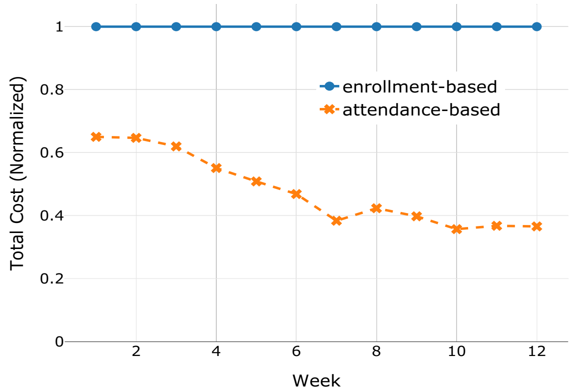

Figure 4: enrollment-based vs. attendance-based cost.

enrollments, and (b) attendance-based, where the room capacity needs to be larger than the number of attendees. These constraints are captured by:

$$o_i \leqslant \sum_{j=1}^{R} C_j x_{i,j} \; \forall \, j \tag{5}$$

where o_i can is either e_i for enrollment-based approach, or a_i for attendance-based approach.

5.2 MILP Optimization and Results

Algorithm: We employ Mixed Integer Linear Programming (MILP) algorithm to solve our problem. MILP is conducive for a problem that has a linear objective function subjected to linear constraints with integer variables. Eq. 2 is used as an input objective function, Eq. 3 defines equality constraints, and Eq. 4 and Eq. 5 defines inequality constraints. The optimization variable $x_{i,j}$ is forced to be binary based on Eq. 3. We assume that there are some spare rooms available for optimal allocation – we use only one spare room of 100 seats capacity. For our optimization problem, we use real data obtained over 8 weeks from courses that are operating in the 9 classrooms.

Results: We run our optimization algorithm for the two approaches (i.e. enrollment-based and attendance-based) for individual one-hour window on each day of a week, and obtain the cost of allocation for each window. We note here that the cost of using a room for an hour is proportional to its capacity. We then compute the weekly cost of allocation by adding hourly costs across the week. We plot the weekly total cost in Fig. 4 – total costs are normalized with respect to the enrollment-based scenario as a baseline (solid blue line). Unsurprisingly, the enrollment-based approach results a constant cost as it tries to meet the fixed constraints every week. On the other hand, total cost obtained from attendance-based allocation (dashed orange lines) falls gradually due to falling pattern of attendance for majority of courses. This suggests that campus can benefit from 52% cost saving over 8 weeks of operation by employing a dynamic allocation of classrooms.

6 CONCLUSION

In this paper we have outlined our experiences in designing and deploying an occupancy sensing system for a real campus environment. We undertook a lab evaluation of various commercial sensors and compared them in terms of cost, ease of operation, and accuracy. We then deployed our beam-counter based system in 9 real classrooms of varying sizes across campus, and collected data over a period of 12 weeks covering 250+ courses, which we release to the public. Our data and visualization reveal interesting insights into course attendance patterns and class utilization measures. Based on this real data, we developed an off-line optimization method for dynamic allocation of courses to classrooms based on attendance rather than enrollments, and showed gains of 52% in room costs.

7 ACKNOWLEDGMENTS

The authors thank the anonymous reviews who gave us valuable feedback on our submission, and a special thanks to our shepherd, Gian Pietro Picco, for his help and suggestions during the revision process.

REFERENCES

[1] 2017. EvolvePlus: Overhead Thermal Counter. https://goo.gl/ugtW1L. (2017).
[2] 2017. EvolvePlus: Wireless People Counters with Remote Data Viewing. https://goo.gl/dAUe8G. (2017).
[3] 2017. Steinel Australia: Presence Control PRO IR Quattro HD. https://goo.gl/T8zvNF. (2017).
[4] 2018. Classroom occupancy data. https://smartcampus.unsw.edu.au/roomoccupancy/data/. (2018).
[5] 2018. Visualization of classroom occupancy. https://smartcampus.unsw.edu.au/roomoccupancy/. (2018).
[6] B. Balaji, J. Xu, A. Nwokafor, R. Gupta, and Y. Agarwal. 2013. Sentinel: Occupancy Based HVAC Actuation using Existing WiFi Infrastructure within Commercial Buildings. In *Proc. ACM SenSys*. Roma, Italy.
[7] E. K. Burke, P. Cowling, J. D. L. Silva, and B. McCollum. 2001. Three Methods to Automate the Space Allocation Process in UK Universities. In *Proc. Springer Practice and Theory of Automated Timetabling III*. Berlin, Heidelberg.
[8] British Council. 2012. *The shape of things to come: higher education global trends and emerging opportunities to 2020.* Technical Report. British Council.
[9] B. Dong, B. Andrews, K. P. Lam, M. Hö Ynck, R. Zhang, Y. Chiou, and D. Benitez. 2010. An information technology enabled sustainability test-bed (ITEST) for occupancy detection through an environmental sensing network. *Energy and Buildings* 42, 7 (July 2010), 1038–1046.
[10] D. Varley E. K. Burke. 1998. Automating Space Allocation in Higher Education. In *Selected Papers from the Second International Conference on Practice and Theory of Automated Timetabling II*. Springer-Verlag, London, UK, 20–36.
[11] E. K. Burke, D. Varley. 1998. Space allocation: An analysis of higher education requirements. In *Practice and Theory of Automated Timetabling II: Second International Conference*. Springer Berlin Heidelberg, Berlin, Heidelberg, 20–33.
[12] G. Fierro, O. Rehmane, A. Krioukov, and D. Culler. 2012. Demo Abstract: Zone-level occupancy counting with existing infrastructure. In *Proc. ACM BuildSys*. Toronto, Ontario, Canada.
[13] S. K. Ghai, L. V. Thanayankizil, D. P. Seetharam, and D. Chakraborty. 2012. Occupancy detection in commercial buildings using opportunistic context sources. In *Proc. IEEE PerCom*. Lugano, Switzerland.
[14] Y. Hou and G. K. H. Pang. 2011. People Counting and Human Detection in a Challenging Situation. *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans* 41, 1 (January 2011), 24–33.
[15] S. S. K. Kwok and E. W. M. Lee. 2011. A study of the importance of occupancy to building cooling load in prediction by intelligent approach. *Energy Conversion and Management* 52, 7 (July 2011), 2555–2564.
[16] K. S. Liu, J. Francis, C. Shelton, and S. Lin. 2017. Long Term Occupancy Estimation in a Commercial Space : An Empirical Study. In *ACM/IEEE IPSN*. Pittsburgh, USA.
[17] K. S. Liu, E. V. Pinto, S. Munir, J. Francis, C. Shelton, M. Bergés, and S. Lin. 2017. COD: A Dataset of Commercial Building Occupancy Traces [Dataset]. (Sept. 2017). https://doi.org/10.5281/zenodo.996587
[18] R. Melfi, B. Rosenblum B. Nordman, and K. Christensen. 2011. Measuring building occupancy using existing network infrastructure. In *2011 International Green Computing Conference and Workshops*. Orlando, FL, USA.
[19] F. Paci, D. Brunelli, and L. Benini. 2014. 0, 1, 2, many – A classroom occupancy monitoring system for smart public buildings. In *Proc. Conference on Design and Architectures for Signal and Image Processing*. Madrid, Spain.
[20] Y. P. Raykov, E. Ozer, G. Dasika, A. Boukouvalas, and M. A. Little. 2016. Predicting room occupancy with a single passive infrared (PIR) sensor through behavior extraction. In *Proc. ACM UbiComp*. Heidelberg, Germany.