An Overview of Occupancy Estimation using Edge ML

Dmitri De Vaz[a], Jayson Bursill[b], R. Christopher Kwong[c], and Gamal Mustapha[d]

[a]P.Eng, Machine Learning Engineer, Delta Controls Inc.
[b]P.Eng, Ph.D, Data Scientist, Delta Controls Inc.
[c]P.Eng, Chief Technology Officer, Delta Controls Inc.
[d]P.Eng, Director of Product Development, Delta Controls Inc.

Abstract

Occupancy estimation is a feature of the O3 Edge and Sense multisensors that uses a suite of non-intrusive sensors and machine learning to estimate the number of occupants in the space. Since the only inputs are environmental sensors, occupancy estimation is a privacy preserving, self-contained, approximate-people-counting solution. The feature is capable of estimating the occupancy in meeting rooms and small office environments with a typical error of less than ± 2 people with a 5 minute reporting interval. In addition, the binary occupancy detection (i.e. unoccupied or occupied) true positive rate is 99%. This paper presents an overview of the occupancy estimation feature including model development, evaluation, application examples, advantages, and limitations.

Keywords — occupancy, machine learning, HVAC, sensors, controls, building automation

1 Introduction

This paper discusses the occupancy estimation feature of the O3 Edge and Sense multisensors [1]. Occupancy estimation uses on-board environmental sensors and a machine learning model to estimate the number of occupants in the space in real-time. We give an overview of the feature and some typical use cases of the technology and discuss it's advantages and limitations. This paper is based on the first experimental release of the occupancy estimation feature. Figure 1 shows the ideal location of the O3 Edge for occupancy estimation.

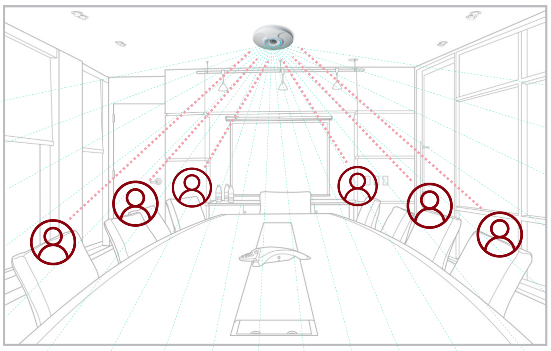

Figure 1: Ideal placement of the O3 Edge in a meeting room.

The effective field of view for the occupancy estimation feature approximately covers a circle with a radius of 3 meters as shown in figure 2.

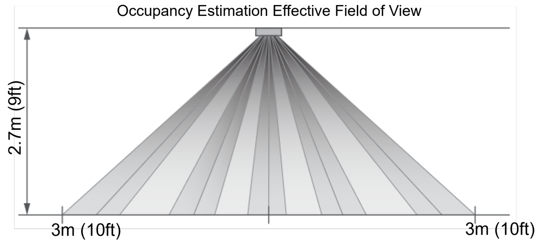

Figure 2: Approximate field of view when mounting at 2.7m (9ft) ceiling height.

The rest of this paper is organized as follows. In section 2 we discuss different levels of existing occupancy systems. Section 3 presents occupancy estimation as it exists in the O3 Edge. Sections 4 and 5 discuss advantages and limitations of occupancy estimation and section 7 presents practical use cases and applications.

2 Background

Occupancy detection systems vary widely in features and capabilities. In this section, we discuss the different levels of existing occupancy solutions.

2.1 Motion Detection

Pyroelectric Infrared (PIR) sensors have been successfully used to provide motion sensing capabilities to thermostats, security systems, and lighting control systems for many years and have become ubiquitous in the Heating, Ventilation, and Air Conditioning (HVAC), lighting, and security industries. They work by using multiple sensing elements that produce a small voltage in response to a warm object moving within the sensor's Field of View (FOV). The signal is then amplified by on-board electronics and compared to a threshold to determine whether there is motion or no motion in the space. PIR sensors typically use a faceted lens called a fresnel lens that helps capture infrared light from all angles and focus it onto the detection elements.

PIR sensors can only detect moving objects hence they are often referred to as motion sensors. They cannot detect objects that remain stationary in the field of view. As a result, PIR sensors by themselves usually do a poor job of detecting binary occupancy, especially when the occupants remain relatively stationary which is common in many real-world scenarios such as meeting rooms.

2.2 Binary Occupancy

Determining whether the space is occupied or unoccupied is referred to as binary occupancy detection. This level of occupancy detection is above motion detection but below occupant counting. For example, the O3 Edge combines PIR sensors, microphones, and IR thermopiles using sensor fusion to determine whether or not the space is occupied. By incorporating information from the PIR sensors, microphones and IR temperature sensors the O3 Edge is able to determine whether or not occupants are in the space even if they are relatively stationary. While this is a significant improvement over motion detection, it would be even more useful if the O3 could determine how many occupants were in the space. Knowing this would enable several important use cases such as Demand Controlled Ventilation (DCV) and space utilization analytics for long-term building owner decision making.

2.3 Occupancy Counting Systems

Occupancy counting systems often use stereo cameras to count the number of occupants in the space. Recent advancements in machine learning and computer vision have drastically increased the accuracy of these systems. The major drawback to these systems is the level of intrusiveness on the occupants. Since cameras are used, the solutions are not privacy preserving. Another drawback is the relative cost of camera based occupancy systems compared to alternatives.

Time-of-Flight (ToF) based solutions are an alternative to camera based systems that use infrared light to provide depth information. ToF based solutions are meant to be placed over doorways and count the inflow and outflow of occupants into and out of a space. These systems can achieve high accuracy and are privacy preserving. There are a few drawbacks for these types of systems. The first is that many rooms have several doorways requiring a device to be installed on each one. Second, some rooms may have large doorways where the narrow FOV of the ToF sensors would not be suitable. Third, these systems tend to require expensive electronics to process the ToF data onboard causing the per-unit-price of these systems to be relatively high. Some solutions opt to send the data out to a cloud server to be processed instead so that the electronics can remain lightweight but these solutions have the drawback that it is not self-contained since occupant information is sent out over the network. Lastly, being a flow based counting system, if a count error occurs it will likely persist until a predefined reset period occurs. This could result in incorrect occupancy for several hours.

Sensor	**Description**
Composite Temperature	The estimated occupant temperature based on air temperature and IR temperature.
Relative Humidity	The occupant height adjusted relative humidity computed from an onboard RH sensor and the composite temperature.
Luminosity	Light intensity measured from an onboard light and color sensor.
PIR Sensor	Standard PIR sensor
Microphones	Two digital MEMs microphones

Table 1: Types of sensors used as input in the occupancy estimation algorithm.

3 Occupancy Estimation

The objective of occupancy estimation is to estimate the number of occupants in the space. Occupancy estimation on the O3 Edge uses a suite of privacy preserving environmental sensors and machine learning to estimate the number of people in the space in real-time. A machine learning model was trained on thousands of hours of sensor data and ground truth occupancy collected from meeting rooms around the world to provide broad variability in the data. It is called *estimation* because the model produces an approximate number instead of the precise number of occupants in the space. This is due to using only privacy preserving environmental sensors instead of a camera based system. In it's current implementation, occupancy estimation on the O3 Edge is designed to work in meeting rooms and small offices with standard geometry. Further details of the quantified limits of the current algorithm are presented later in this paper. Occupancy estimates are generated once per 5 minute interval and are available over BACnet and MQTT. Please refer to the application guides for more information on how to access the estimated occupancy [2].

3.1 Method Overview

The O3 Edge acquires environmental sensor data from infrared and air temperature sensors, a relative humidity sensor, a PIR sensor, a luminosity sensor, and two microphones as shown in table 1. The sensor data is fed as inputs to a machine learning model. The model's purpose is to take these time series signals as input and process it to extract the number of occupants in the space.

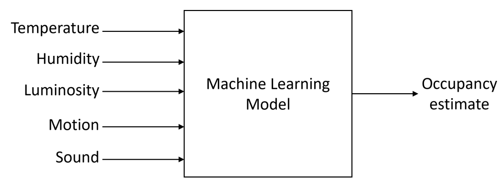

Figure 3: Inputs and outputs of the machine learning model.

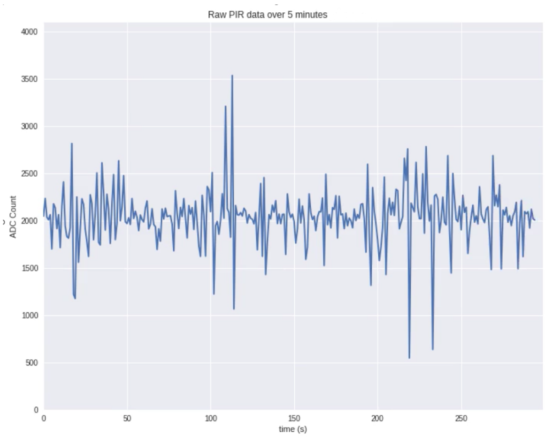

Figure 4: Raw PIR data collected over 5 minutes, used as one of the inputs to the occupancy estimation system.

Intuitively, we should expect temperature to rise in proportion to the number of occupants in the space. In practice, however, it is not straightforward. In a typical meeting room, the HVAC system actively cools or heats the space depending on the room temperature and it's deviation from a predetermined set-point. This means that the temperature in the room is not only influenced by the number of occupants but also by the specific properties of the room's temperature control system and related hardware. Another example is the PIR sensor that has a wide field of view and can often *see* all occupants in the space. We are interested in the aggregate behaviour of all of the occupants in the FOV. However, the resulting PIR signal is a superposition of the aggregate behaviour of the occupants (i.e. the component of interest), large oscillations due to individual occupants which obscures the aggregate information, and signal noise (usually caused by the large amplification stage required to boost the small PIR signal). The objective of the machine learning model is to extract the parts of these signals that are well correlated with the number of occupants in the space while discarding the parts that are not. A sample of one of the input sensors, the PIR sensor, is shown in figure 4 and a block diagram of the high level steps of the model is shown in figure 6.

3.2 Model Development

The model was developed by training on many hours of data collected from real meetings. Over 2000 hours of sensor data and the corresponding true occupancy count (ground truth occupancy) was collected from several meeting rooms around the world.

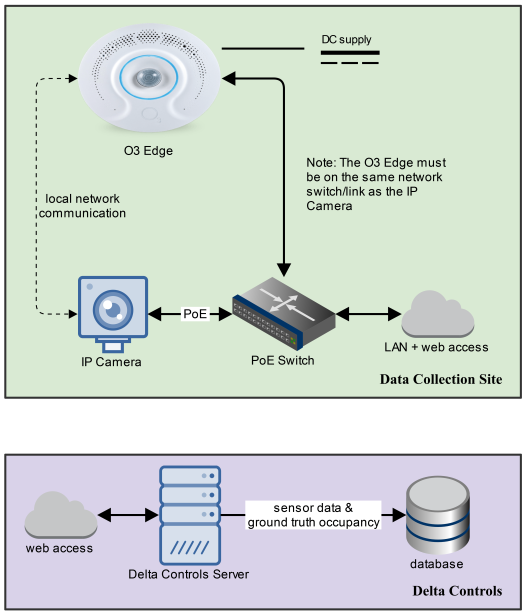

Figure 5: Data collection kit that collects anonymous sensor data along with ground truth occupancy data

A data-collection kit was sent out to participating Delta Controls partners in order to facilitate the collection of anonymous environmental sensor data from the O3 Edge alongside ground truth occupancy in the form of still images taken from an I.P. camera. The sensor data and the images are compressed and securely transmitted to Delta Controls' servers using TLS encryption. The data-collection kit contains an O3 Edge and associated power supply, an IP camera, a network switch, and a setup guide.

3.3 Edge ML

Edge ML refers to efficient machine learning algorithms that can run on severely resource-constrained edge IoT devices ranging from small 8-bit microcontrollers to larger linux-based embedded systems. Today, millions of microcontrollers are used in several industries such as predictive maintenance, agriculture, wearables, smart buildings, and healthcare. Edge ML can enable new abilities or enhances existing abilities for these devices. In the case of the O3 Edge, it enables a highly sought after feature that isn't possible to achieve without using ML, a privacy preserving, self-contained occupancy estimation system. The fact that the system is self-contained and works completely

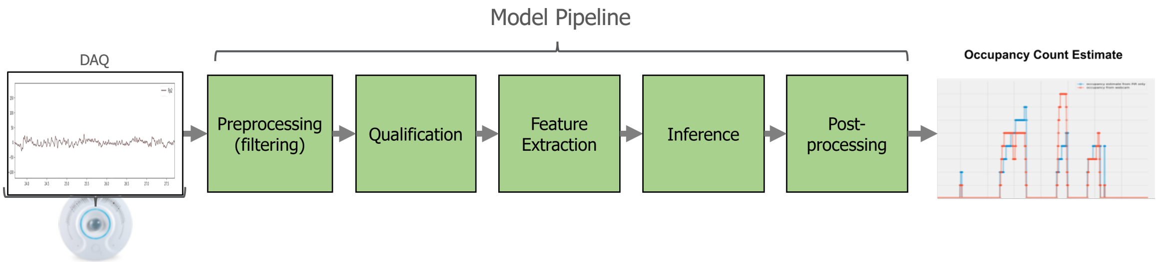

Figure 6: High-level steps from raw sensor data to occupancy estimate.

offline without an internet connection is important for privacy, security, ease of deployment, and energy usage.

The model running on the O3 Edge was trained on powerful servers using industry standard tools for data science and machine learning model development (python data science ecosystem). Since these tools produce models that are most often run in a powerful desktop or server environment (often with high end graphics processing units) the resulting model worked well but was much too computationally demanding to deploy to a resource-constrained embedded system such as the O3 Edge directly. For practical purposes, the model needs to be able to work alongside all of the other tasks that the O3 Edge is doing such as handling BACnet objects, computing edge analytics, running Node-RED and executing GCL code to name a few examples. Therefore, there was a considerable amount of effort that went into the quantization, model pruning, and model optimization in order to reduce the amount of storage, memory, and computational resources required to run the model directly on the O3 Edge without sacrificing model accuracy.

3.4 Evaluation

To evaluate the model we needed to assess it's performance on *new* meetings. The model was evaluated on a test dataset which consisted of several hours of meeting room data from environments that the model was not trained on. This allows us to assess how well the model generalizes to new spaces that it has not encountered in training, that is, the model's generalization performance. Good generalization performance is extremely important for having a practical solution since the model would not be as useful if it required extensive training for each new environment. To evaluate the model we need to quantify it's performance relative to the ground truth occupancy over all of the meetings in the test dataset. For each meeting, we compute two performance metrics that describe the extent of correctness of the model relative to the true occupancy over the course of a meeting. The first is the percent of time where the error is within two people (PW2) as shown in equation 1. That is, the percent of time where the model's estimate of the number of occupants is within 2 occupants of the true value. This metric gives an intuitive sense of how well the model performs. For example, a PW2 of 90% means that the model's estimates are within ±2 of the true value 90% of the time.

$$PW2 = \frac{n_{error<2}}{n_{total}} \tag{1}$$

While the PW2 metric is useful for quantifying the performance of the model in an intuitive way, one shortcoming is that it does not capture the performance when the model is not within ±2 of the true occupancy. This is handled by the second metric, the mean absolute error (MAE). It is the average absolute error between the model's estimate and the true occupancy count. The MAE quantifies how *wrong* the model is in general and how much error we should expect on average. For example, an MAE of 2 people suggest that the model's occupancy estimates are accurate to within 2 people, on average. An MAE of 1.8 people, suggests that the model error is between is between 1 and 2 people on average but is usually closer to 2 people. MAE is shown in equation 2 where n is the number of samples, y_i is the true value and x_i is the model's output.

$$MAE = \frac{1}{n}\sum_{i=1}^{n}|y_i - x_i| \tag{2}$$

In addition to evaluating the accuracy of the model during the course of a meeting, it is just as important to evaluate the ability to detect the start and end of the meetings themselves.

We treat this as a binary occupancy problem (occupied or unoccupied) and evaluate the system by quantifying the false positive (reporting occupancy when the room was unoccupied) and false negative (reporting unoccupied when the room was occupied) rates.

The current version of the O3 Edge's occupancy estimation feature is typically accurate to within ±2 occupants. The current experimental model has a PW2 of 87% and an MAE of 1.5 people. Figures 7 and 8 show typical meetings with 6 occupants and 9 occupants, respectively.

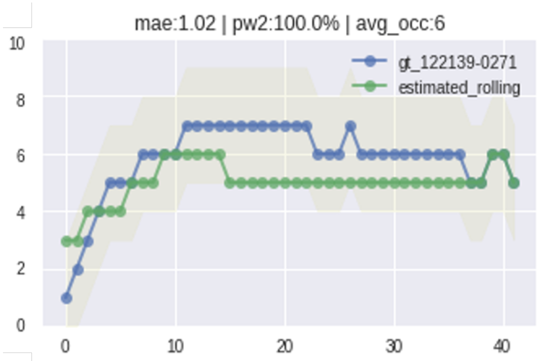

Figure 7: Occupancy performance for a meeting with 6 occupants. Blue: true occupancy, Green: estimated occupancy, Yellow band: ±2 region centered around the true occupancy. Vertical axis: number of occupants, horizontal axis: meeting duration (minutes)

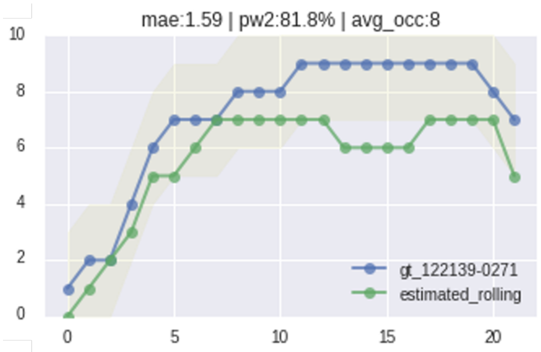

Figure 8: Occupancy estimation during the course of a meeting with 9 occupants. Blue: true occupancy, Green: estimated occupancy. Yellow band: ±2 region centered around the true occupancy. Vertical axis: number of occupants, horizontal axis: meeting duration (minutes)

The binary occupancy detection accuracy is 99%. The full confusion matrix is shown in table 2.

		predicted	
		unoccupied	occupied
actual	unoccupied	93%	7%
	occupied	**1%**	99%

Table 2: Binary occupancy confusion matrix showing the performance of the model for detecting the occupied vs unoccupied states.

Note that the 7% false positive rate is solely due to time-alignment and not isolated false positives. For example, when occupants leave the room after a meeting the occupancy estimation system may stay in the occupied state for another 5 minutes before going to the unoccupied state. This is a result of the system summarizing the occupancy in the previous 5 minutes. This type of false positive is in contrast to one that is generated while in the unoccupied state, causing an erroneous unoccupied to occupied transition.

4 Advantages

In this section we discuss several advantages of the occupancy estimation feature as it exists in the O3 Edge.

Privacy: Occupant privacy is preserved since only anonymous data are used as inputs to the machine learning model. For example the PIR sensor data shown in figure 4.

Cost: Occupancy estimation on the O3 Edge is a relatively low cost solution at less than $400 CAD per unit. In contrast, camera, radar, and ToF based solutions are all considerably more expensive. For example, Density's ToF based solution starts at $895 USD per unit and requires an annual subscription of $795 USD [3]. Vivotek's stereo camera based solution retails at $3999 USD for two units [4].

Self-Contained: The occupancy estimation feature is completely self-contained. It does not require an internet connection and does not use any cloud connectivity to estimate the number of occupants in the space. Furthermore, it only uses locally acquired data from the local building automation system (BAS).

Fast: Occupancy estimates are generated once every five minutes. This fast reporting interval enables real-time applications such as DCV. Importantly, the reaction time to large, intermittent occupancy fluctuations is much faster than CO_2 based systems. This is discussed in detail in section 7.3.

Energy Efficient: In general, occupancy estimation is an energy efficient solution. It uses less energy compared to camera, ToF, and Radar based occupancy systems. In addition, edge computing is initiated after initial occupancy is detected. This minimizes overall processing cycles.

Zero Calibration: Currently, the occupancy estimation feature requires no calibration or re-calibration process. The only thing that is required is to place the O3 Edge centrally in the meeting room or office. The device will automatically produce occupancy estimates every 5 minutes as well as keep a historical record of hourly occupancy summaries for the last 30 days.

5 Limitations

In this section, the limitations of the occupancy estimation feature are presented.

5.1 Limited Environments

The occupancy estimation feature currently only works in standard meeting rooms with area up to 20x20 ft and small office environments with standard geometry. Supported ceiling height range is from 8-12 ft. The device must be centrally mounted in the occupant space with no visible obstructions that would block the O3's field of view. Due to limitations of the privacy preserving environmental sensors used as input to the model, occupancy estimation will not work well in highly dynamic environments such as fitness centers or hallways. Any deviation from the above requirements may produce erroneous or inaccurate data.

5.2 Precision

During the start and end of a meeting, when occupants enter or leave the space there may be a momentary spike in the estimated occupancy due to the increased activity as compared to normal activity during a meeting. If a momentary spike occurs, it will diminish shortly and settle back closer to the true number of occupants within a few minutes.

Occupancy estimation will not be as accurate as a camera based people counting system since it is using non-intrusive, low cost, environmental sensors to estimate occupancy. We expect that the precision of the occupancy estimation feature will be sufficient for most real-world use-cases but understand that there exists other scenarios where more precise systems are required.

6 Comparison to other Occupant Counting Technologies

Table 3 compares different occupancy systems across several different criteria. All factors except *Self-contained* and *Cost* are graded on a 4 point scale with 4 being the best possible score and 1 being the lowest possible score. *Accuracy* refers to the people counting system's typical accuracy relative to the other solutions, *Privacy* refers to the amount of intrusiveness in the solution. *Self-contained* determines whether or not the occupancy can be computed without data leaving the place where it originated. For example, a solution that requires sending data to a cloud server would not be considered self-contained. *Response time* is based on the typical reporting interval of the occupancy system with 4 points representing a fast response time. *Versatility* refers to the number of different environments and applications where the system can work well. Finally, *cost* refers to the relative cost of each of the occupancy systems.

7 Use Cases and Applications

There are many practical uses for the occupancy estimation feature. In this section we present some use cases for occupancy estimation: building analytics for long-term decision making, real-time information for dynamic room allocation, and demand-based ventilation. In the demand based ventilation section, we highlight some of the drawbacks of the existing

Occupancy System	Accuracy	Privacy	Self-contained	Response time	Versatility	Cost
Cameras	●●●●	●○○○	no	●●●●	●●○○	high
Beak-Break IR	●○○○	●●●●	no	●○○○	●●●○	low
MAC Address Tracking	●○○○	●○○○	no	●●○○	●●○○	medium
Seat Sensors	●○○○	●●●○	yes	●●○○	●○○○	low
Thermal Camera	●●●○	●●●○	yes	●●○○	●●●○	medium
Ultrasonic	●○○○	●●●●	no	●●○○	●●○○	high
Time of Flight + AI	●●●○	●●●○	no	●●●●	●●●●	high
Radar + AI	●●○○	●●●●	no	●●●●	●●●○	high
O3 Sensors + AI	●●○○	●●●●	**yes**	●●●●	●●○○	**low**

Table 3: Types of occupancy systems

methods and present how occupancy estimation can be used instead.

7.1 Building Analytics

Building analytics are data analytics extracted from smart building technologies such as IoT devices, networked sensors, and building management systems. Building analytics help building managers make important long term decisions. Occupancy estimation provides information about which rooms are over or under utilized. In particular, given the occupancy estimates for each room over time $occ_{est}(t)$ and the number of samples, n, the room's average occupancy over a given period can be computed as:

$$\mu_{occ} = \frac{1}{n} \sum_{\Delta T} occ_{est}(t) \tag{3}$$

Where ΔT is the time period of interest. Occupancy utilization can then be calculated over the same time period as a percentage:

$$utilization = \frac{\mu_{occ}}{occ_{max}} \times 100 \tag{4}$$

Where occ_{max} is the maximum designed occupancy for the room. To facilitate these computations, the O3 Edge's edge analytics module stores 30 days of hourly mean estimated occupancy (μ_{occ}) as well as the maximum and minimum estimated occupancy. This allows computing the hourly utilization easily by dividing the provided μ_{occ} by the room's maximum occupancy, occ_{max}.

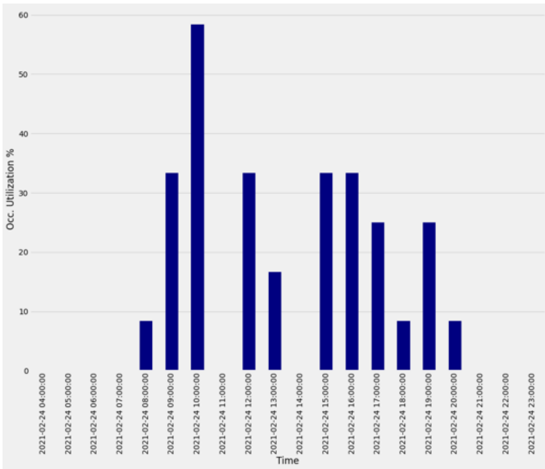

Figure 9: Hourly occupancy data from the Edge-Analytics module on the O3 Edge

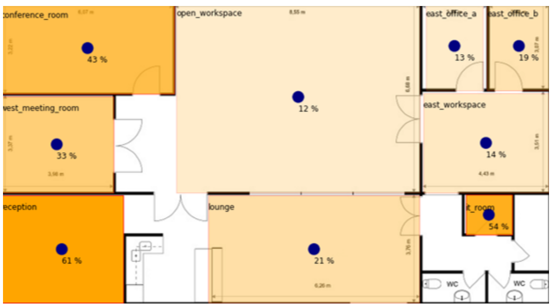

Figure 10: Visualizing occupancy utilization via a heatmap

A useful way to visualize occupancy utilization data is in the form of a heat-map as shown in figure 10. This allows building managers to see which rooms are overused or underused and make adjustments accordingly. For example, deciding to restrict access to infrequently used meeting rooms to save heating/cooling and lighting energy. Coordinating cleaning services, performance bench-marking, and forecasting are some other applications of building occupancy analytics.

7.2 Real Time Information

Since a new occupancy estimate is computed every 5 minutes, real-time applications are possible. For example, the feature can be used to maintain room capacity limitations. This is typically accomplished by using a graphical display shown on the outside of the meeting room or on a central monitor where all meeting rooms are displayed.

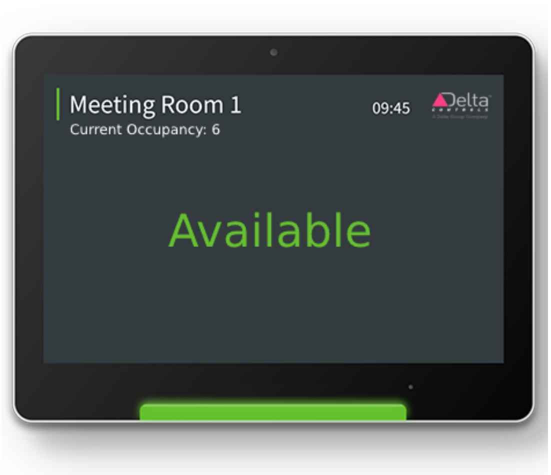

Figure 11: Meeting room digital sign showing availability and current number of occupants.

If the room occupancy is within the maximum capacity then the sign will show that the room is available as in figure 11. A real-time heat-map is also possible that would show which meeting rooms are available at a glance. Using real-time occupancy information can increase meeting room allocation while improving the experience of the occupants.

7.3 Demand-based Ventilation

Demand based ventilation can have a dramatic impact on both energy savings and occupant comfort. HVAC systems are known to be the highest contributors of building-related energy consumption and consume 48% of the total energy consumption in buildings [5]. Typically, HVAC systems are based on the maximum design occupancy of the rooms. This often results in increased HVAC energy consumption [6]. On the other hand, reducing the ventilation has been shown to negatively affect occupant comfort [7]. Poor ventilation can also result in the increased spread of diseases and viruses such as COVID-19. CO_2 sensors have been used successfully to implement DCV however they suffer from several problems. CO_2 sensors have been used to coarsely estimate the number of occupants in the space. While CO_2 can be effective for occupancy detection, CO_2 spreads very slowly in the environment and takes a substantial amount of time to build up to a level that registers on a single sensor in the space. This causes the resulting approximate occupancy based on the CO_2 signal to have a significant time delay. In addition, sensor placement and passive ventilation can significantly affect CO_2 levels. For example, Pantelic et al. shows that occupants have a personal CO_2 cloud and if the CO_2 sensor is not placed near the occupant's CO2 cloud the readings could be delayed or incorrect [8]. Therefore the placement of these sensors is extremely important and often not considered when used in real-world applications. Another problem with CO_2 sensors is that they require periodic re-calibration with external gas to maintain accuracy since automatic baseline calibration relies on assumptions that don't hold true in many environments.

The American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) guideline 62.1-2019 [9] states that the amount of outdoor air that should be diffused into the space is directly proportional to the number of occupants as shown in equation 5.

$$R = (N_{occ} \times 7.5cfm) + (A_{sqft} \times 0.6cfm) \quad (5)$$

Where R is the outdoor air flow rate in cubic feet per minute (cfm), N_{occ} is the number of occupants in the space, and A_{sqft} is the room area in square feet. Instead of using CO_2 as a proxy for the number of occupants, it would be more effective to estimate the number of occupants directly. In addition to ASHRAE, the International Energy Conservation Code now requires the engineer to use DCV strategies in any space that is at a density equal to or greater than 25 people per 1000 square feet, with few exceptions [10]. DCV is known to be a good energy conservation strategy, in general, for any space that is expected to be intermittently occupied, regardless of whether the code requires it.

Instead, DCV uses building occupancy level information to dynamically adjust air ventilation based on the amount of occupancy in the building. This can significantly reduce a building's energy consumption. Research suggests that having access to occupancy counts and patterns, building automation systems (BAS) can more efficiently control energy usage and occupant comfort, resulting in up to 80% reduction in HVAC-related energy consumption [11]. Switching to a DCV based system is highly effective even if the occupancy counting system is basic. For example, a study of 81 buildings in Norway showed that using a binary occupancy based DCV system reduced the energy usage by an average of 49% [12].

We can use the occupancy estimates generated from the O3 Edge directly as N_{occ} in equation 5 to compute the required ventilation rate. However we often want to intentionally introduce *some* time lag in the form of a filter in order to smooth out the control signals sent to the air handler. A first order exponentially weighted moving average filter (EWMA) works well and is straightforward to implement as shown in equation 6. By using the parameter α we can introduce as much smoothing (and associated lag) as desired for the particular room and air handler conditions.

$$s_t = \alpha x_t + (1 - \alpha) s_{t-1} \tag{6}$$

Where s_t is the smoothed occupancy estimate at current time step, s_{t-1} is the smoothed occupancy estimate for the previous time step, x_t is the occupancy estimate for the current time step, and α is a parameter that controls the amount of smoothing where $0 < \alpha < 1$. Values of α close to 1 have little smoothing whereas values near 0 have a large amount of smoothing. To initialize the filter, we can set it to the value of the first occupancy estimate, i.e. $s_0 = x_0$. By using this filter, the engineer can control the time-dynamics to accommodate different room parameters and HVAC equipment.

8 Conclusion and Future Work

In this paper we presented an overview of the occupancy estimation feature of the O3 Edge and discussed a few applications: DCV, real-time occupancy, and long term decision making based on edge analytics. We presented advantages and limitations of the feature and compared the solution with existing people-counting solutions on the market. The first version of occupancy estimation is able to achieve a typical accuracy of ± 2 people in meeting room environments. We are working on expanding our data collection efforts to cover more scenarios, edge cases, and environments in order to further improve the precision and reliability of the occupancy estimation feature and also to expand it's potential use cases. Occupancy estimation serves as a low cost, privacy preserving, self-contained, approximate people counting solution that has sufficient precision for many practical applications.

References

[1] Delta Controls Inc. *O3 Sense and Edge Multisensors.* URL: `https://deltacontrols.com/products/o3/`.

[2] Delta Controls Inc. *O3 Sense and Edge Multisensors Support.* URL: `https://support.deltacontrols.com/Products/O3Edge`.

[3] Forbes. *Density, Platform For Counting People Inside Buildings, Raises $51 Million In Series C.* URL: `https://www.forbes.com/sites/igorbosilkovski/2020/07/28/density-platform-for-counting-people-inside-buildings-raises-51-million-in-series-c`.

[4] Network Camera Store. *Vivotek CCS-2D People Counting Crowd Control Kit.* URL: https://www.networkcamerastore.com/vivotek-ccs-2d-people-counting-crowd-control-kit-2-door.

[5] Nasruddin et al. "Optimization of HVAC system energy consumption in a building using artificial neural network and multi-objective genetic algorithm". In: *Elsevier* Sustainable Energy Technologies and Assessments.35 (2019), pp. 48–57.

[6] C. Wang, K. Pattawi, and H. Lee. "Energy saving impact of occupancy-driven thermostat for residential buildings". In: *Elsevier B.V.* Energy and Buildings (2020), p. 211.

[7] Z. Yang et al. "Systematic approach to occupancy modeling in ambient sensor-rich buildings". In: *Simulation* Sustainable Energy Technologies and Assessments.90(8) (2014), pp. 960–977.

[8] J. Pantelic et al. "Personal CO2 cloud: laboratory measurements of metabolic CO2 inhalation zone concentration and dispersion in a typical office desk setting". In: *Journal of Exposure Science and Environmental Epidemiology* Sustainable Energy Technologies and Assessments.30(2) (2020), pp. 328–337.

[9] ASHRAE. *ASHRAE Standard 62.1-2019.* URL: https://ashrae.iwrapper.com/ASHRAE_PREVIEW_ONLY_STANDARDS/STD_62.1_2019.

[10] The International Energy Conservation Code. *Chapter 4: Commercial Energy Efficiency, SECTION C403.2.6.1.* URL: https://www3.iccsafe.org/cs/committeeArea/pdf_file/EC_15_71_17.pdf.

[11] J. Brooks et al. "Energy-efficient control of under-actuated HVAC zones in commercial buildings". In: *Elsevier B.V.* Energy and Buildings.93 (2015), pp. 160–168.

[12] M. Mysen et al. "Occupancy density and benefits of demand-controlled ventilation in Norwegian primary schools". In: *Energy Build* Sustainable Energy Technologies and Assessments.37(12) (2005), pp. 1234–1240.