A Novel Occupancy Detection Solution using Low-Power IR-FPA Based Wireless Occupancy Sensor

Aravind K. Mikkilineni[a], Jin Dong[b,*], Teja Kuruganti[c], David Fugate[d]

[a] *Imaging, Signals, & Machine Learning Group, Oak Ridge National Laboratory, One Bethel Valley Road, Oak Ridge, TN USA 37831*
[b] *Building Envelope & Urban Systems Research Group, Oak Ridge National Laboratory, One Bethel Valley Road, Oak Ridge, TN USA 37831*
[c] *Modeling & Simulation Group, Oak Ridge National Laboratory, One Bethel Valley Road, Oak Ridge, TN USA 37831*
[d] *Systems Analysis, Aerojet Rocketdyne, 555 Discovery Drive, Huntsville, AL, 35806*

Abstract

Significant energy savings can be achieved by operating heating, ventilation and air conditioning control systems with indoor occupancy measurement information. This paper presents a novel plug-and-play occupancy sensing method which will enable the temporal minimization of building energy consumption to meet building usage behavior without privacy concerns. The proposed wireless occupancy sensing platform is based on long-wave infra-red (LWIR) focal-plane arrays (FPAs), or thermal imagers, that detect thermal energy rather than visible light. We developed an advanced sensor package consisting of multiple thermal imagers with low-cost optical enhancements to increase field of view and increase sensitivity to occupant detection (filtering building clutter). These imagers can be coupled with radio frequency and ultrasonic-based radar to enhance data collection at key occupant zone boundaries to improve accuracy. Standard filtering and estimation techniques from the image processing and computer vision communities are introduced to overcome the accuracy issues suffered by traditional PIR based sensing, especially when occupants remain relatively still. Accurate low-level counting of individuals can be achieved with minimal impact on privacy. The proposed occupancy detection method can

*Corresponding author
Email address: dongj@ornl.gov (Jin Dong)

work as a retrofit to enable real-time understanding of the building operational state and usage behavior. Sensor data obtained from real test building zones was used to test the efficacy of the method.

Keywords: Long-wave infrared (LWIR) sensor, focal plane array (FPA), Human movement detection, Occupancy detection, Occupancy counting, HVAC control, Sensor data.

1. Introduction

1.1. Background

Heating, ventilation, and air-conditioning (HVAC) systems consume 30% of building energy and comprise 50% of building electricity consumption. Most modern buildings continue to control room temperature and ventilation assuming maximum occupancy rather than actual occupancy. High energy consumption is largely due to conventional HVAC and lighting systems using open-loop control strategies that cannot dynamically respond to changes in occupancy [1]. Recent studies [2, 3, 4, 5] show that there exists a significant potential for energy savings by temporally matching building energy consumption and building usage via occupancy-based control. With more than 134 million houses in the United States [6], this presents a tremendous opportunity to decrease energy consumption and reduce inefficiencies. Occupancy-based control (OBC) can achieve up to 30% energy savings by temporally matching building energy consumption and building usage [7, 8, 9].

The goal of this paper is to design a high-performance occupancy sensor node and optimize installation locations based on building layout in order to achieve accurate occupancy detection and minimize installation and energy costs. Commonly used devices for detecting occupancy in buildings include passive infrared and ultrasonic sensing technologies typically used for lighting or security applications. In recent years, alternative technologies for indoor occupancy sensing have been studied, including microwave radar, acoustic, light barriers, video cameras, biometric systems, pressure pads, and electric field sensors. Many of these have been used to detect occupancy for safety or security with little application specifically toward building control due to integration cost and challenges with respect to detection accuracy and precision.

1.2. Related Work on Occupancy Sensing

Significant work has been done to investigate the potential of improved detection accuracy for real-time occupancy-based HVAC control. A comparison of the different occupancy systems is given in Table 1 [7]. As an observation from the table, passive infrared (PIR) sensors are relatively cheap compared to video camera systems.

Type of sensor	Resolution	Number of occupants	Person identification	Person localisation	Initial cost
PIR	Low	No	No	No	Low
Ultrasonic	Low	No	No	No	Low
Microwave	Low	No	No	No	Low
Sound	Low	No	No	No	Low
Light barriers	Low	Yes	No	No	Low
Video	Very high	Yes	Yes	Yes	High
Biometric	High	Yes	Yes	No	High
Pressure	Low	No	No	No	Medium

Table 1: Comparison of occupancy sensing technologies [7].

Based on the occupancy detection technique, there can be several types of sensing systems. Common occupancy detection systems typically used in office buildings for demand-driven applications include direct sensors such as carbon dioxide (CO_2), PIR, ultrasonic, image, pressure pads, radars and electric field sensors; or implicit sensors such as energy consumption, phone, Wi-Fi and computer activity detection sensors [10].

Moreover, a fusion of multiple sensors has also been studied in occupancy detection for demand-driven applications [11, 12]. Previous research [13, 12, 14] has introduced a system for estimating building occupancy based on data from a sensor network including PIR, video cameras, and CO_2 detectors.

It is noted that concentration of gases in a space in parts-per-million (PPM). Hence, the CO_2 detector becomes a commonly used tool for the measurement of occupancy in buildings for demand-driven control of HVAC systems because the amount of CO_2 in a space can provide an estimate of user presence as well as count [15, 16, 17]. However, its application in occupancy estimation is hampered by the relatively poor real-time performance and measurement precision due to slow gas mixture [18, 13].

1.2.1. Vision-based Occupancy Detection

Vision-based systems which rely on camera images and video analysis techniques are often used in buildings for security purposes, though their use has also been explored for occupancy measurement in buildings. In general, computer vision-based processing for human tracking includes following stages: modeling of environments, detection of motions, classification of moving objects, tracking, understanding and description of behavior, human identification, and the fusion of data from multiple cameras [19, 20]. However, such systems can raise serious concerns in the following aspects:

1. It is sensitive to ambient illumination, which changes during the day and at night [21];
2. The computational/communication load for continuous visual surveillance is heavy;
3. Using cameras may introduce privacy concerns, causing an uncomfortable feeling or even an adverse psychological impact for the residents due to the fact they are being observed [22].

Most other sensing methods involving sound, ultrasonic, electric field and all the implicit ones suffer serious problems in accuracy, privacy and false ON/OFF [11, 23, 24, 25, 10].

1.2.2. PIR-based Occupancy Detection

In order to overcome these weaknesses of video-based detection, some researchers employ PIR sensors [26, 27, 28, 29, 30, 31]. PIR sensors are non-intrusive and are only sensitive to changes in infrared radiation such as that caused by human motion. This makes them robust to interference caused by cluttered background and illumination variance [32]. PIR sensors have been widely employed for human detection systems due to its low cost and power consumption, small form factor, and unobtrusive and privacy-preserving interaction. A multi-modal system consisting of a PIR sensor and a regular camera, which distinguishes entry/exit motions and ordinary body movements, is proposed in [33].

The PIR sensor is the most commonly used technology for occupancy sensing in buildings, especially for lighting control [34]. Even though a solid body of work on occupancy sensing in building environments has been conducted in the past decades [35, 36, 37, 10, 38], there remain technical challenges with PIR sensors that hinder the development of innovative solutions for occupancy sensing. It is well known that the application of PIR sensing in buildings is limited due to several major flaws (as claimed in [10, 39]): (a) PIR sensors can only provides coarse binary information (occupied or not) [40, 20, 11]; (b) PIR sensors require a direct line of sight between the sensor and occupants in a space; (c) PIR sensors require continuous motion to function effectively (i.e. occupants are seating still or the sensor's view is impeded by other objects) [41]; (d) PIR sensors can be triggered by alternative thermal currents from hot coffee/tea, 3D printers, HVAC systems [20] or pets [39].

These limitations of the existing PIR-based occupancy detection systems may lead to incorrect control actions and degraded human comfort since building HVAC control differs significantly from instantaneous lighting control. In particular, ramping room temperature up or down is not instantaneous owing to inertia in the underlying mechanical systems (e.g., compressors) and heat transfer. Although an occupancy monitoring system can provide estimates in near real time, if the estimates are used in a reactive manner, then the HVAC control will likely leave occupants uncomfortable until target temperatures are met. To ensure occupant comfort, robust predictors are needed to adaptively condition a given space [42, 43].

It is worth mentioning that most of the aforementioned systems do not provide zone level occupancy information and were not designed or tested for tracking multiple mobile occupants [20]. Tracking mobile occupants is essential in achieving better occupancy estimation accuracy (since historical information of the same occupant is used) and control efficiency (since the occupancy movement should be accounted for when adjusting the HVAC operations).

1.3. Motivation and Research Objective

Given this background, we pursue in the present study a long-wave infrared (LWIR) focal-plane array (FPA) based occupancy detection platform to address the major concerns presented by existing solutions. We therefore use thermal imagery, which captures infrared radiation instead of visible light, and creates an image whose pixel values represent temperature [44]. This is a hybrid method that combines the benefits of both vision-based and PIR-based techniques. Compared with the traditional camera-based measurement, people can not be identified in these thermal images, eliminating potential privacy issues. A positive side effect of thermal imaging is that detection/tracking can often be reduced to a trivial problem by utilizing the existing large body of algorithms for image processing and computer vision. Hence it enables more sophisticated occupancy estimation/tracking than traditional PIR sensor measurement.

LWIR sensors have been studied in the vehicle society for occupant- and driver-posture analysis as well as pedestrian detection [45, 46]. In the buildings' envelope area, LWIR thermography can be employed to evaluate building materials [47], detect construction defects [48], determine the heat losses in buildings [49], building energy diagnostics [50], and other problems relating to humidity [51, 52]. However, to the best of our knowledge, there exist few studies in the literature using LWIR sensors for building occupancy detection and/or tracking.

The complexity of the proposed algorithms along with memory and energy consumption must be balanced such that the complete solution can be carried out in an embedded architecture. The developed system must therefore be an optimal trade-off between occupancy detection accuracy and algorithmic complexity.

In particular, the overall objectives are:

1. To detect the presence of a person in its environment without being disturbed by the presence of animals or other moving objects.

2. The sensor must be robust to changes in illumination, ambient heat sources, and be able to protect users' privacy.
3. The complexity of the proposed algorithms should be considered with limited computational and energy resources, allowing the system to function in an embedded architecture without external power.

To address the challenge of minimizing the required number of sensors for any given building, a graph-based optimal sensor placement algorithm is proposed in this paper. In addition to this optimal design in placement, we have also developed our own embedded board integrating LWIR and PIR sensors with wireless capabilities to provide energy saving occupancy monitoring and estimation services.

There is a desire with any sensing system to achieve perfect accuracy in the detections and/or decisions that it makes. However, it is important to also realize that the quest for perfection can come at a cost of additional algorithm complexity, increased sensor cost, and greater energy consumption at the sensor node itself. These tradeoffs need to be considered in the larger picture in combination with the building control system with which the sensor will eventually be integrated. A typical building control system will have a minimum control resolution which may correspond to more than one person. Therefore any sensing accuracy/resolution that is finer than that of the control system provides no additional benefit toward total building energy reduction.

For example, the sensor node described here has been shown to perform well for detection and counting of humans entering and leaving different zones within a building. In its current instantiation it may not be able to differentiate a cat, or perhaps a large dog, from a human. In an office environment this is perhaps not as significant an issue as in a home environment. However, the fact that an FPA sensor is used can allow additional algorithmic improvements to be made to potentially address that issue. Additionally, sensor fusion methods making use of multiple sensors on the platform along with machine learning could be investigated as a way to disambiguate various thermal bodies in the field of view of the sensor.

The rest of the paper is organized as follows. In the remainder of the Introduction, we review related literature on technologies that underlie the considered occupancy sensor solutions and developments. The proposed wireless sensor-driven occupancy detection system provides a complete infrastructure to realize sensor placement as well as development. We present the architecture of the occupancy detection system in Section 2 to realize the system solutions and services. A technical realization to realize the services is then described in the following three sections. In Section 3, we consider optimal placement of occupancy sensors to cover the entire area of interest. A solution for offering accurate low-level counting of individuals is described in Section 4. Services for utilizing the developed occupancy detection algorithm and analytics using collected sensor data are described in Section 5. Conclusions and discussions are provided in Section 6.

2. System Architecture

A straightforward approach currently used to add smart OBC to a building is to simply place a sensor in every room and hallway. However, this approach is expensive, requiring many more sensors than are potentially necessary. Indeed, the largest gains in building efficiency can typically be achieved not at the individual room level, but at the zone level. Sensors necessary to provide sufficient information for adaptive control at the zone level can reduce installation cost significantly while still allowing measurement of system state and detection of abnormalities. An optimal distribution of sensing nodes will also consider the lowest-cost sensing option for each specific node necessary to meet the required accuracy goals. There are several architectures that the sensing system could follow with respect to integration into the building. Centralized architectures use significant energy for communication from node to central location, but minimal energy in computation at node level. Distributed architectures expend less energy on communication and increased energy for node-level data processing.

In designing the sensing nodes and the resulting system architecture, it is important to consider the physical distribution of the nodes along with the local sensing requirements at each node. Some nodes may require a simple PIR and temperature sensor (e.g., a discrete entrance), whereas others may require a thermal imager (e.g., large cafeteria). Irrespective of the type of low-level sensor on each node, the core information that each node obtains from the sensors would ideally be the same. Architectures for deployment, data processing, and control hierarchy will drive the cost-effectiveness of the solution.

2.1. Wireless Sensor Node

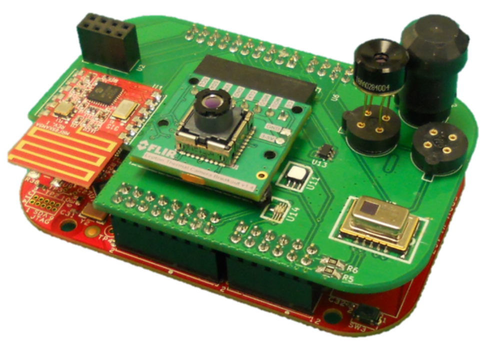

Figure 1: Wireless sensor node for advanced occupancy sensing and control.

To help foster a similar culture and capability for occupancy sensing a common sensing platform is developed with which researchers can test and validate distributed multi-modal sensing algorithms. The sensor node that has been designed and fabricated as part of this project is shown in Figure 1. The sensor node includes several LWIR sensors (thermal imagers), a PIR sensor, as well as sensors for temperature, relative humidity, and ambient light. A sub-GHz radio is also included to allow communication back to a central control unit to update occupancy count. As the capabilities and performance of the sensing node progresses, a more abstract interface will be developed though which various sensors and algorithms can be easily integrated into the node with minimal necessary low-level hardware or software knowledge.

Each node can be powered on in one of two modes: (1) Control, and (2) Sense. The desired mode is chosen by pressing one of the buttons on the node while powering it up.

Control nodes can plug into a USB port on a computer, providing a serial interface through which to allow wireless control of multiple sense nodes. Control actions that are currently defined include sensor selection, sensor orientation, begin occupancy sensing, and querying of occupancy count. Additional actions are being added to complete a minimal functional set for testing.

Sensor nodes are expected to run headless, however if powered through USB connected to a computer then debugging and other information about its operational state can be obtained through its serial interface.

2.2. Sensors

The highest resolution sensor being considered is the FLIR Lepton which is an 80x60 microbolometer based imager available with either a 50 or 25 degree field of view and a maximum frame rate of 8.6 frames per second (FPS). This sensor has the highest power consumption out of all the sensors described drawing approximately 150 mW. One particularly interesting feature with this sensor is that in addition to outputting data as a single 8-bit or 14-bit value for each pixel representing radiated temperature, it can also automatically apply colorization to the sensed temperature data and output the data in an RGB format which can greatly assist in later processing tasks such as clustering that would be used for detection and tracking of occupants using standard image processing methods.

In addition to the bolometer based Lepton sensor, two thermopile based devices are available, the Melexis MLX90621 which features a 16x4 pixel array, and the Panasonic AMG883X which features an 8x8 pixel array. These sensors are much lower resolution compared to the Lepton, however in the case of the Melexis offerings are available in much wider and narrower fields of view which may be useful depending on the local building geometry where the sensor node will be installed. In addition, the non-square aspect ratio of the Melexis sensors means that they are highly orientation dependent.

To enable comparisons between various sensors and to allow for multi-modal sensing as the project progresses, all of the sensors listed in Table 2 were selected to be integrated into the node Figure 2. Current cost estimates are also included, however it should be noted that only a subset of these sensors would be used in an actual commercial node offering and an industrial manufacturer could typically receive even lower quantity pricing.

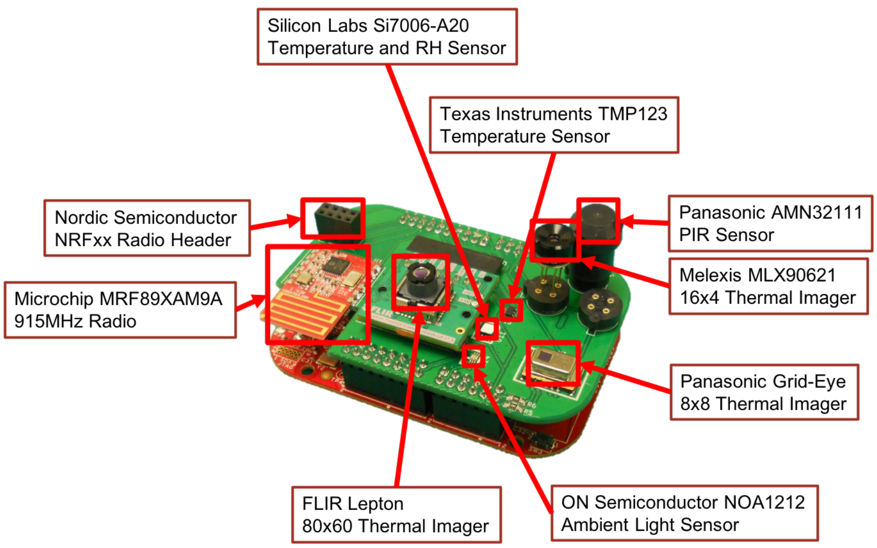

Figure 2: Close look at the sensor node.

2.3. Node Architecture

The node must utilize wireless communication and be low-power being able to operate on an internal battery for a reasonable length of time, or permanently by using one or more energy harvesting methods.

The board that was selected for construction of the sensor node is the Freescale/NXP FRDM-KL46Z evaluation board. This board is built around a Freescale/NXP KL46Z256 micro-controller which contains a low-power ARM

Table 2: Selected sensors for sensor node implementation

Type	Manufacturer	Model	Array size	Cost
LWIR	FLIR	Lepton	80×60	$175
LWIR	Melexis	MLX90621	16×4	$52
LWIR	Panasonic	AMG88	8×8	$16
PIR	Panasonic	AMN32111	NA	$11
Temperature + RH	Silicon Labs	Si7006-A20	NA	$1
Temperature	Texas Instruments	TMP!23	NA	$1
Ambient Light	ON Semiconductor	NOA1212	NA	$0.21

Cortex-M0+ processing core with 256 kB of program memory, 32 kB of RAM, and a maximum clock rate of 48MHz. Included on the board are a built-in debug interface, an ambient light sensor, and a LCD display that can be used to display the node status while in operation. Although the micro-controller on this board is potentially larger (with respect to memory) and more powerful than absolutely required, it does allow rapid algorithm development after which algorithm efficiency and code minimization can be pursued to allow a smaller less powerful, and cheaper, micro-controller to be used in the final node design if required to meet cost or other design goals.

The Freescale/NXP Kinetis micro-controllers have been selected for several reasons. These micro-controllers are some of the lowest cost ARM based micro-controllers available on the market and are available in a wide array of physical sizes and packages that allow better matching to the specific set of sensors and board layout eventually used. In addition, the Freescale development tools are one of the few that do not restrict code size and other capabilities without a paid license. This is especially important when considering that the algorithms and source code for the developed platform will be shared with collaborators and will need to be accessible to other parties for testing, verification, and integration without undue difficulty. Another reason to select these specific evaluation boards is that they are supported by the ARM Mbed platform and operating system which is designed to simplify the development of devices for the Internet of Things (IoT) and enable code reuse across ARM devices. The complete node is packed with a 3D-printed case (as shown in Figure 3), which is locally manufactured at Oak Ridge National Laboratory (ORNL)'s Manufactuing Demonstration Facility.

The first step in creation of the sensing node is to determine an appropriate mapping between each sensor and the selected development boards available interfaces. A connectivity diagram of the selected sensors and how they interface with the FRDM-KL46Z development platform is shown in Figure 4. The development platform is the block in the center, and all the blocks surrounding it are the sensors that are externally interfaced to it.

3. Sensor Placement

To determine the ideal or optimal placement of sensors for determining the system state, it is important to understand what the underlying causes of the state changes are and how they evolve the system state. Observation of state changes, such as a person passing between zones or exiting the building, provides information on the current state. As advanced estimation methods are developed, the goal is to be able to operate with a number of sensors at or below the number needed for full state capture.

The proposed innovation will be threefold:

- First, we provide efficient approximation algorithms that select a small number of sensors to optimize the estimation error based on real occupancy across different zones.
- Second, we develop an optimal deployment algorithm for the occupancy sensors (the number of sensors is the optimized result from the first accomplishment).

(a) (b)

(c)

Figure 3: Packed node inside 3D-printed cases

- Third, we solve the scheduling problem to improve the lifetime of such a sensor network based on the deployment scheme.

Any building can be modeled as an un-directed graph depicting zone connectedness within the building and the external environment. Consider an arbitrary floor plan for an office building. An automated graph model can be directly generated based on the connectedness of each zone. Figure 5 shows a graph model for a floor plan with four zones. Each vertex of the graph represents a zone, and every connection between zones is represented by an edge. The outer gold boundary is the boundary of the system delineating between the interior

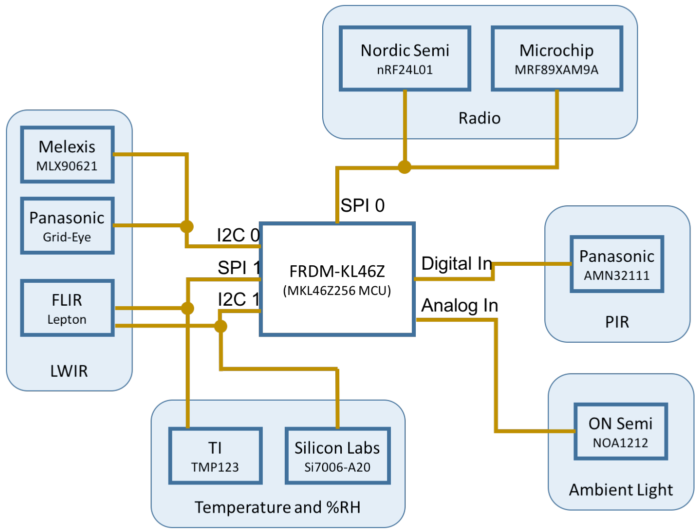

Figure 4: Connectivity diagram of sensor node.

of the building and the exterior. Each zone is depicted as a node in the graph, shown as Z_i. It should be mentioned that it is possible for multiple edges to connect each zone such as illustrated between Zone 4 and the exterior of the system.

Each edge represents a boundary between zones that people can cross, such as a door or a point along a hallway. Note that the aggregate system occupancy is determined solely by information available on the green edges. Zone i occupancy, Z_i, is solely determined by information available on those edges with a vertex on Z_i, which includes both inter-zone boundaries and system boundaries. The total number of people in the system is then the sum of all the individual zone counts $\|Z\|_1$ or simply $\|Z\|$ from hereon.

Note that the entire state of the system is given by the zone count vector Z and that the system state is only affected by movement of people across the edges of the graph. In particular, the system occupancy $\|Z\|$ is only affected by movement across the system boundary (green edges), whereas each individual zone occupancy Z_i is affected by movement across inter-zone boundaries (blue edges) as well as system boundaries. Given these observations it is arguable that

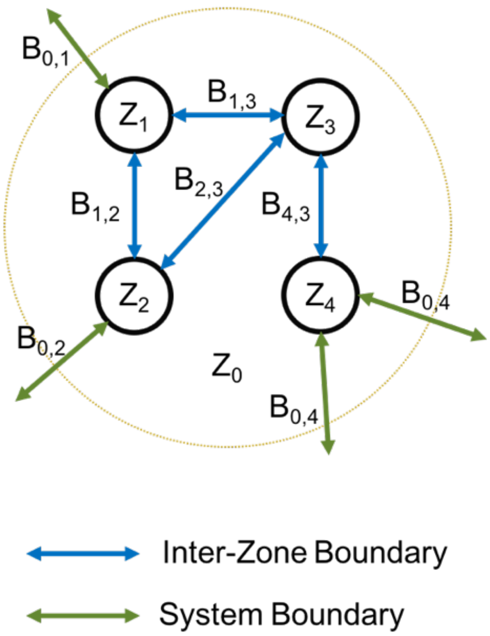

Figure 5: Sensor Placement.

information available through observation of the zone boundaries is sufficient to completely determine the system state.

Each boundary is typically representative of a physical boundary in a building such as a doorway, or a virtual boundary such as a specific location such as where two hallways meet or join together. It can be assumed that there exists a small area within each zone from which each connected edge or boundary can be observed. The majority of any given zone potentially provides no information about edge behavior since the boundaries may not be observable such as from inside an office or closet. Figure 6 shows a closeup view of Zone 2 from Figure 5. The dashed lines delineate the areas within the zone from which the associated edges are observable. The system boundary in this case is observable from a small area, and the two inter-zone boundaries are both observable from another small area.

It can be argued that within a zone, one sensor per boundary is sufficient to

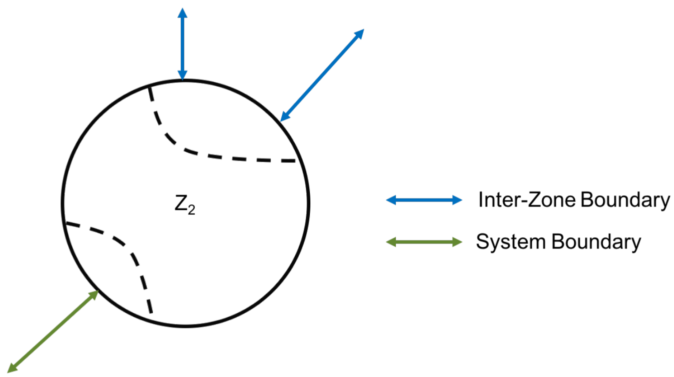

Figure 6: Closeup view of Zone 2.

capture occupancy count for that zone. Minimizing the number of sensors necessary to provide zone count information is therefore inherently upper bounded by the number of edges at a given vertex (3 in this case). In this case since the two inter-zone boundaries are potentially both visible using a single sensor reducing the sensor count by 1 (down to 2 sensors). Reducing the sensor count further in this case requires estimating occupancy based on data from a single sensor somewhere within the zone. Based on this proposed placement algorithm, we will discuss two use cases, one for office environments and another for warehouse environments.

3.1. Edge Coverage for Office

Considering the limited geometrical size of normal offices and the detection range of PIR sensors, we can directly deploy single sensor for each concerned boundary as shown in Figure 7. We use the same floor plan as mentioned earlier in Figure 5, where four interconnected zones are involved. As shown in the figure, all the potential edges and door entrances are monitored by sensors denoted by blue S_1, S_2, S_3, S_4. Here we assume the field of view of each sensor is larger than 90 degrees, allowing the use of as few as four sensors to cover the entire office area. Several observations are noted below:

1. By tuning the direction of Sensor 1 (S_1), it is potentially able to observe both Door 1 and the zone boundary between Zone 3 and Zone 1;
2. Though there are two paths connecting Zone 3 and Zone 4, its possible to use Sensor 2 (S_2) alone to monitor the edge as long as the detection range is long enough;
3. Sensor 3 (S_3), with an appropriate field of view, can view the boundary between Zone 2 and Zone 3, and also Door 2 which crosses the system boundary;
4. Sensor 4 (S_4) will cover all the possible paths from Zone 1 to Zone 2 including the hallway and a door to one of the offices in Zone 2.

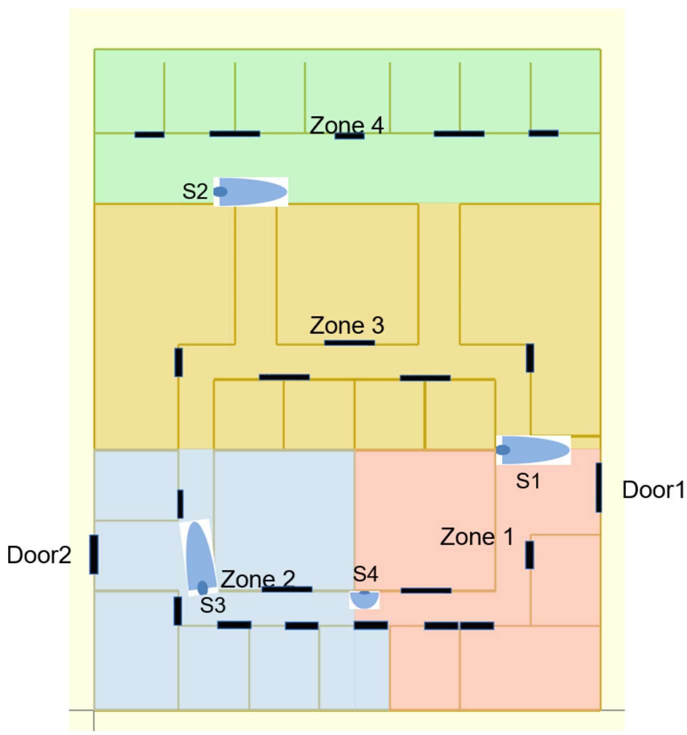

Figure 7: Example sensor deployment in an office.

It should be remarked that the aforementioned sensor placement is based on the assumption that the detection range and field of view meet the corresponding requirements, respectively. Additional sensors may be necessary if these condition are not satisfied. For example, if S_2 is not able detect movement at the far hallway, then a separate senor should be placed close to that hallway. Or if S_3 is unable to view both Door 2 and the hallway transition simultaneously, then two vertically oriented sensors may be considered instead of one single sensor.

3.2. Edge Coverage for Warehouse

Now, we will study the second example about deployment in a warehouse. A top view of a typical warehouse floor plan is depicted in Figure 8. The warehouse is divided into 6 zones denoted with different colors. There are several shelving units placed in each zone. The warehouse has two main entrances at the bottom, and 5 emergency exits around the other three sides of the building. Notice that the dotted red line connecting Ex 1 and Ex 5 (in Figure 8), as well as the vertical zone boundaries, can reach over 100 ft in length.

For simplicity, we focus on the red dotted line in Figure 8. It can be easily generalized to any boundary/edge between zones. After abstraction, the red dotted line can be viewed as the horizontal black line in Figure 9. Several pairs of sensors can be placed along the line to cover the entire boundary. To achieve the best coverage performance, two sensors are deployed side by side towards opposite directions. Particularly, the red ones face toward the left, while the blue ones to the right in Figure 9.

4. Occupancy Counting Algorithm

Image and video processing is well studied and many methods exist for object detection, segmentation, and tracking.Accurate low-level counting of individuals can be achieved through standard particle or Kalman filtering or particle flow methods to eliminate certain errors obtained with other methodologies using a similar combination of sensors. Depending on the geometry of a building, it is also possible for an imager to sense a much larger area, allowing much

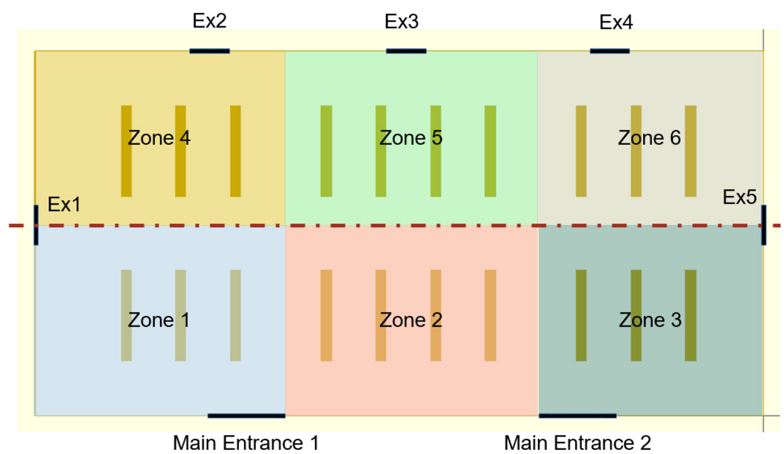

Figure 8: Example sensor deployment in a warehouse.

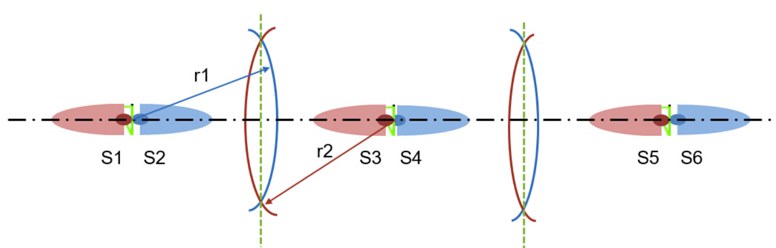

Figure 9: Sensor coverage of the long edge.

higher-order occupancy statistics to be obtained. Additionally, a 3-dimensional thermal profile of a building can be achieved through a low-complexity template matching method to obtain these pieces of information from the thermal image at the sensor node itself with little computation. A volumetric thermal profile of the building can be constructed in real time through sensor data fusion that can be used to identify register outflow temperature, leaks around doors or windows, and stratification. The combination sensor can be realized in a small form factor (< 1 in.^{2}) at low cost.

4.1. Control Logic

Appropriate software control of each node is necessary to meet the various requirements including low power operation and multiple sensor support. A flowchart outlining the basic ideal operation of a sensing node is illustrated in Figure 10. When the node first powers up, it will begin by initializing all the peripherals within the MCU itself which includes the various communication buses necessary to communicate with each sensor. The node will then register itself on the sensing network indicating to the control node that it is active and present, and afterward place the radio in a standby mode until it is needed again.

All of the Tier 1 sensors will then be enabled and the MCU will be placed in a low-power sleep state. At this point the only significant power consumption will be by the Tier 1 sensors including PIR, temperature, humidity, and illumination. Each of these sensors can be configured to trigger an interrupt when a certain threshold is met waking up the MCU. The thresholds at which each sensor triggers can be selected based on motion or other state changes (such as temperature) in the sensing region that may be indicative of an impending change in occupancy.

Once the MCU awakens it will immediately enable the Tier 2 sensor(s), in this case one or more LWIR FPAs. The data from the FPA is processed frame by frame. A frame is captured from the FPA and detection of occupants is performed. If no occupants are detected and no activity has been recorded by either the Tier 1 or Tier 2 sensors for some predefined period of time then the Tier 2 sensors are disabled and the MCU goes back to the sleep state to await further activity. If instead occupants are detected, then each occupant is individually tracked. If the result of the tracking between the current and previous frame indicates a change in zone count then an appropriate message is sent to the control node indicating the change. Otherwise another frame is captured from the FPA and processing continues again at the detection stage.

4.2. Occupancy Detection Algorithm

The algorithm for detection and tracking is based on basic image processing methods including thresholding, region growing, and blob detection. These methods have existed since at least the 1960s and are well covered in any introductory image processing textbook. The application of these algorithms for the goals of this project are described below for the case where the sensor is mounted above a walkway looking straight down.

The acquired sensor data can be viewed as an image in which each pixel's value is a temperature of some object in the scene. The objects of interest in this case are people who are assumed to be at a higher temperature than the surrounding objects in an office environment. Given this observation the first step in the processing chain is to threshold the image by setting any pixel value below some threshold temperature to 0, and all other pixels to 1.

One way to do this would be to select a fixed threshold temperature just below human body temperature. However, a fixed threshold will not work reliably in this case due to the fact that radiated body temperature is in general lower than internal body temperature and may vary significantly between individuals. The temperature distribution in a scene without any humans present may also vary between pixels and over time either triggering false detections when no humans are present, or by missing detections due to a human appearing below the selected temperature threshold due to circumstances such as just walking indoors on a cold day.

Instead of a fixed threshold, a per-pixel adaptive threshold is used. The threshold at each pixel in this case is set as 2 °C above a rolling average over a fixed number of frames. In our implementation we have chosen to average over 20 frames, and only those frames with no detected blobs (background only) are used in the calculation. Setting the threshold slightly above the average allows small variations in the background temperature without triggering a false detection. The assumption here is that the majority of the variation is from the sensor itself, and that the background is only slowly varying in time.

After thresholding, a blob detection algorithm is used to find all the contiguous regions in the image that were identified as non-background. As part of the blob detection, each blob is assigned a unique label and statistics are obtained including its centroid and size. As the detected blobs move through the scene, they cross several virtual boundaries that are used to determine whether a blob needs to be actively tracked, as well as whether the blob has contributed to a change in occupancy.

For the doorway installation scenario the occupancy count would change as individuals walk through the doorway. Assuming that the doorway lies vertically in the center of the image, that means occupancy changes as individuals move between the left and right halves of the image. To assist in the detection of these transitions the image is split as shown in Figure 11. The zone boundary in this case lies vertically through the center of the image. The two zones are labeled Zone 1 and Zone 2 on the left and right sides of the zone boundary respectively. Both zones are separated by a transition region that straddles the physical zone boundary.

This is done for two reasons. The first is that the transition sub-region defines a region of interest (ROI) outside of which any detected blobs can be ignored. Any blob that crosses the doorway threshold must pass through the ROI, and restricting the tracking algorithm to only this region reduces computation time. The second reason for using two boundaries is that it introduces hysteresis into the tracking, allowing a blob to potentially cross the doorway threshold multiple times without affecting the occupancy count until that blob fully enters or exits into one of the adjacent zones. An example of this would be a person standing in the doorway itself while talking to a colleague.

Any blob that is detected within the ROI is tracked. This means that its position when initially found is stored along with its most current position. The zone from which the blob entered the ROI is assumed to be the zone that it is closest to upon entry. When a tracked blob exits the ROI through either boundary it is removed from the list of tracked blobs and the occupancy count is appropriately updated. Any movement of the blob within the ROI itself does not affect the occupancy count.

Blob tracking is performed using the blob statistics mentioned earlier. Those statistics allow determination to be made as to whether a blob in the current frame overlaps with a blob in the previous frame. If any overlap exists it is assumed that both are the same blob which has simply moved between frames. If the matching blob in the previous frame is being tracked, then its current location is updated. If there is no matching blob and it is in the ROI in the current frame, then it is added to the list of actively tracked blobs until it exits the ROI.

As the blob centroid moves across some defined boundary in the image the zone count is either incremented or decremented appropriately.

4.3. On-board Communication Solution

For delivering a complete plug-and-play occupancy detection solution, each occupancy sensor node can be selected as control or sense mode during the power-on process. As shown in Figure 12, occupancy measurement collected from the three grey-colored sensors will be transmitted back to the control red-colored node using wireless communication. The control node can be connected to any computer or energy management systems via USB. This hierarchical structure enables the sensor platform to be conveniently deployed at any location without relying on other communication infrastructure.

This dedicated communication capability has several advantages over the WiFi-dependant solution.

- Simplify the platform setup since it avoids requirement and procedure of getting the WiFi connection ready.
- Power conserving since the wireless communication is only activated when there is an occupancy count change as described in the main sensing logic (Figure 10).
- Cyber-resilient since this will be separated from the main network which

is fragile to outside intruder. This helps protect potentially sensitive occupancy information.

5. Real-field Test and Results

This section presents results for validating the hardware and software designs including the occupancy counting performance, accuracy, communication and calculation capabilities. To cover different scenarios of human crossing zone boundaries, we tested people enter/leave the zone separately as well as simultaneously.

It is worth mentioning that there are no existing practical solutions for multiple people tracking and counting since the commercial ready ones (laser beam based) can count one person at a time, while the camera based solutions have serious concerns related with privacy. Therefore, we want to show our occupancy detection node fills this gap in state of the art.

Testing occured in a typical office building at ORNL. The developed occupancy sensor node was installed on the top jamb of a doorway (7 ft high) looking directly down at the floor and covering an approximate region size of 16 sq.ft. (at floor level). The node was powered by a USB battery bank and connected to a control node using on-board radio communication as discussed in Section 4.3.

As mentioned before, to minimize communication and centralized computation, our design distributes all the detection, tracking, and occupancy counting to the occupancy sensor nodes. Various requirements including low power operation and multiple sensor support can be met following the control logic defined in Section 4.1. It should be mentioned that we only implemented it in a office hallway which can be considered as a typical zone boundary. This works as a basic proof of concept in that it can count the number of people entering/leaving the zone by sensing only the boundary itself. For more general deployment with a wider doorway or larger area, similar results are expected to be achievable by using higher resolution sensors with an appropriately selected field of view.

A novel blob tracking algorithm is applied based on adaptive thresholds for accurate and reliable detection and tracking of human movement. The images in the first rows of Figures 13 - 15 show the raw sensor output as one or more individuals cross the observed zone boundary. This data uses only the Panasonic Grid-Eye 8x8 pixel thermal imager to show what can be achieved with a minimal number of pixels on target. Any blob that is detected within the ROI is tracked and compared with historical data log until they completely leave the transition region. This enables us to handle the case of more than one object entering the scene at the same time, which is an open question for most of the existing occupancy detection solutions.

In particular, Figure 13 represents a simple case, where only one human crosses the boundary. While Figures 14 and 15 show the scenario when there are two persons present in the scene. Testing results from Figures 14 and 15 show that our proposed algorithms can deal with the complicated scenarios where multiple people crossing the boundary at the same time. In Case 1, one person first entered the scene from the left, then a second people entered from the right. Our occupancy detection algorithm correctly reported two occupants during the test. Similarly, in Case 2, we switched the order of two persons entering the scene. Our algorithm has also succeeded in reporting correct number of occupancy in this case.

The data acquisition frequency used while obtaining these results was approximately 10 frames/sec. It should be mentioned that the sensor was asleep by default, only waking up when a person approached within view of the 170 degree PIR sensor contained within it, at which point a quick rolling multi-frame background estimation was performed and continually updated until either a target entered view or a timeout period with no motion (LWIR or PIR) was reached.

It is worth mentioning that the neither the raw nor processed images in Figures 13 - 15 involve any personally identifying information, which meets our design requirement for privacy protection.

6. Conclusion and Future Directions

In this work we have presented the design and architecture of a novel plug-and-play occupancy sensor that enables temporal minimization of building energy consumption to meet building usage behavior without privacy concerns. We developed an advanced sensor package consisting of multi-pixel thermal imagers with low-cost optical enhancements to increase field of view and increase sensitivity to occupant detection (filtering building clutter). Furthermore, standard filtering and estimation techniques from the image processing and computer vision communities were introduced to overcome the accuracy issues suffered by traditional PIR based sensing, especially when occupants remain relatively still. Accurate low-level counting of individuals has been demonstrated while maintaining minimal impact on privacy. In addition to the occupancy sensor node, we have also introduced a practical graph-based optimal sensor placement algorithm to minimize number and optimize location of the occupancy sensors.

The developed occupancy detection method advances the state-of-the-art by overcoming multiple challenges including accuracy, cost, energy efficiency, and privacy preservation, which makes it suitable as a retrofit device to enable real-time understanding of the building operational state and usage behavior. Contrary to existing methods, the proposed method is immediately operational (plug-and-play) as it does not require time-consuming gathering of detailed information about the physical conditions of the room or the need to wait for extensive training data prior to reliable operation. This will greatly contribute to bridge the gap between occupancy detection and occupancy-based control, especially the emerging model predictive control (MPC). Moreover, our solution will help evaluate the flexibility and possibility of using buildings for grid-efficient interactive purposes.

One immediate step in making this occupancy sensor more popular would be to work with other partners and organizations to test as many of these specific mechanisms as possible, to confirm that they work as expected and evaluate any outstanding issues. The content of this report could then be expanded and converted into a more comprehensive guide on how to customize and enhance the platform for various scenarios. This would provide a rich resource for building owners and researchers who would like to make use of occupancy data. Finally, there is a need to disseminate case studies, data, and results of this work to the building energy efficiency community to share the opportunities described in this study. Many building owners and operators are unaware of the value of and technical methods to gather these occupancy data.

Acknowledgment

The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes. The Department of Energy will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (http://energy.gov/downloads/doe-public-access-plan).

References

[1] J. Zou, Q. Zhao, W. Yang, F. Wang, Occupancy detection in the office by analyzing surveillance videos and its application to building energy conservation, Energy and Buildings 152 (2017) 385–398.

[2] F. Oldewurtel, D. Sturzenegger, M. Morari, Importance of occupancy information for building climate control, Applied energy 101 (2013) 521–532.

[3] S. Goyal, H. A. Ingley, P. Barooah, Occupancy-based zone-climate control for energy-efficient buildings: Complexity vs. performance, Applied Energy 106 (2013) 209–221.

[4] J. Dong, C. Winstead, S. M. Djouadi, J. J. Nutaro, T. Kuruganti, Stochastic modeling of short-term occupancy for energy efficient buildings, in: 4th

International High Performance Buildings Conference at Purdue, 2016, pp. 1–10.

[5] J. Dong, C. Winstead, S. M. Djouadi, J. J. Nutaro, T. Kuruganti, D. L. Fugate, Short-term occupancy prediction for demand based HVAC control, Applied Energy (under review).

[6] U. C. Bureau., State county quickfacts: Usa.
URL https://factfinder.census.gov/faces/nav/jsf/pages/community_facts.xhtml

[7] J. Zhang, G. Liu, R. Lutes, M. R. Brambley, Energy savings for occupancy-based control (obc) of variable-air-volume (VAV) systems, PNNL-22072 (2013) 1–79.

[8] S. Bengea, F. Borrelli, A. Kelman, S. Killough, T. Kuruganti, Y. Lin, W. Manges, S. Narayanan, R. Taylor, D. Vrabie, et al., A wireless platform for energy-efficient building control retrofits, Final report draft, ESTCP Project EW-0938.

[9] S. Goyal, P. Barooah, T. Middelkoop, Experimental study of occupancy-based control of HVAC zones, Applied Energy 140 (2015) 75–84.

[10] T. Labeodan, W. Zeiler, G. Boxem, Y. Zhao, Occupancy measurement in commercial office buildings for demand-driven control applicationsa survey and detection system evaluation, Energy and Buildings 93 (2015) 303–314.

[11] X. Guo, D. Tiller, G. Henze, C. Waters, The performance of occupancy-based lighting control systems: A review, Lighting Research & Technology 42 (4) (2010) 415–431.

[12] B. Dong, B. Andrews, K. P. Lam, M. Höynck, R. Zhang, Y.-S. Chiou, D. Benitez, An information technology enabled sustainability test-bed (itest) for occupancy detection through an environmental sensing network, Energy and Buildings 42 (7) (2010) 1038–1046.

[13] S. Meyn, A. Surana, Y. Lin, S. M. Oggianu, S. Narayanan, T. A. Frewen, A sensor-utility-network method for estimation of occupancy in buildings, in: Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference. CDC/CCC 2009. Proceedings of the 48th IEEE Conference on, IEEE, 2009, pp. 1494–1500.

[14] V. Garg, N. Bansal, Smart occupancy sensors to reduce energy consumption, Energy and Buildings 32 (1) (2000) 81–87.

[15] N. Nassif, A robust co 2-based demand-controlled ventilation control strategy for multi-zone HVAC systems, Energy and buildings 45 (2012) 72–81.

[16] S. A. Mumma, Transient occupancy ventilation by monitoring co2, ASHRAE IAQ Applications (2004) 21–23.

[17] G. Ansanay-Alex, Estimating occupancy using indoor carbon dioxide concentrations only in an office building: a method and qualitative assessment, in: Proc. REHVA World Cong. Energy Efficient, Smart and Healthy Build.(CLIMA), 2013, pp. 1–8.

[18] S. Wang, J. Burnett, H. Chong, Experimental validation of co2-based occupancy detection for demand-controlled ventilation, Indoor and Built Environment 8 (6) (1999) 377–391.

[19] W. Hu, T. Tan, L. Wang, S. Maybank, A survey on visual surveillance of object motion and behaviors, Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on 34 (3) (2004) 334–352.

[20] N. Li, G. Calis, B. Becerik-Gerber, Measuring and monitoring occupancy with an rfid based system for demand-driven HVAC operations, Automation in construction 24 (2012) 89–99.

[21] J. Han, B. Bhanu, Human activity recognition in thermal infrared imagery, in: Computer Vision and Pattern Recognition-Workshops, 2005. CVPR Workshops. IEEE Computer Society Conference on, IEEE, 2005, pp. 17–17.

[22] B. Jansen, F. Temmermans, R. Deklerck, 3d human pose recognition for home monitoring of elderly, in: Engineering in Medicine and Biology Society, 2007. EMBS 2007. 29th Annual International Conference of the IEEE, IEEE, 2007, pp. 4049–4051.

[23] C. Martani, D. Lee, P. Robinson, R. Britter, C. Ratti, Enernet: Studying the dynamic relationship between building occupancy and energy consumption, Energy and Buildings 47 (2012) 584–591.

[24] K. Christensen, R. Melfi, B. Nordman, B. Rosenblum, R. Viera, Using existing network infrastructure to estimate building occupancy and control plugged-in devices in user workspaces, International Journal of Communication Networks and Distributed Systems 12 (1) (2014) 4–29.

[25] M. Milenkovic, O. Amft, Recognizing energy-related activities using sensors commonly installed in office buildings, Procedia Computer Science 19 (2013) 669–677.

[26] Q. Hao, D. J. Brady, B. D. Guenther, J. B. Burchett, M. Shankar, S. Feller, Human tracking with wireless distributed pyroelectric sensors, Sensors Journal, IEEE 6 (6) (2006) 1683–1696.

[27] R. C. Luo, J. Yi-Huei, O. Chen, Robotics human tracking system through wireless pyroelectric sensor system, in: Advanced robotics and Its Social Impacts, 2008. ARSO 2008. IEEE Workshop on, IEEE, 2008, pp. 1–6.

[28] T. Hosokawa, M. Kudo, H. Nonaka, J. Toyama, Soft authentication using an infrared ceiling sensor network, Pattern Analysis and Applications 12 (3) (2009) 237–249.

[29] J. Yun, M.-H. Song, Detecting direction of movement using pyroelectric infrared sensors, Sensors Journal, IEEE 14 (5) (2014) 1482–1489.

[30] Q. Hao, F. Hu, J. Lu, Distributed multiple human tracking with wireless binary pyroelectric infrared (pir) sensor networks, in: Sensors, 2010 IEEE, IEEE, 2010, pp. 946–950.

[31] P. Zappi, E. Farella, L. Benini, Tracking motion direction and distance with pyroelectric ir sensors, Sensors Journal, IEEE 10 (9) (2010) 1486–1494.

[32] X. Luo, H. Tan, Q. Guan, T. Liu, H. H. Zhuo, B. Shen, Abnormal activity detection using pyroelectric infrared sensors, Sensors 16 (6) (2016) 822.

[33] F. Erden, A. Z. Alkar, A. E. Cetin, A robust system for counting people using an infrared sensor and a camera, Infrared Physics & Technology 72 (2015) 127–134.

[34] D. T. Delaney, G. M. O'Hare, A. G. Ruzzelli, Evaluation of energy-efficiency in lighting systems using sensor networks, in: Proceedings of the First ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings, ACM, 2009, pp. 61–66.

[35] Y. Agarwal, B. Balaji, S. Dutta, R. K. Gupta, T. Weng, Duty-cycling buildings aggressively: The next frontier in HVAC control, in: Information Processing in Sensor Networks (IPSN), 2011 10th International Conference on, IEEE, 2011, pp. 246–257.

[36] V. L. Erickson, M. Á. Carreira-Perpiñán, A. E. Cerpa, Occupancy modeling and prediction for building energy management, ACM Transactions on Sensor Networks (TOSN) 10 (3) (2014) 42.

[37] B. Balaji, J. Xu, A. Nwokafor, R. Gupta, Y. Agarwal, Sentinel: occupancy based HVAC actuation using existing wifi infrastructure within commercial buildings, in: Proceedings of the 11th ACM Conference on Embedded Networked Sensor Systems, ACM, 2013, p. 17.

[38] T. Labeodan, K. Aduda, W. Zeiler, F. Hoving, Experimental evaluation of the performance of chair sensors in an office space for occupancy detection and occupancy-driven control, Energy and Buildings 111 (2016) 195–206.

[39] Y. Benezeth, H. Laurent, B. Emile, C. Rosenberger, Towards a sensor for detecting human presence and characterizing activity, Energy and Buildings 43 (2) (2011) 305–314.

[40] H. Zou, H. Jiang, J. Yang, L. Xie, C. Spanos, Non-intrusive occupancy sensing in commercial buildings, Energy and Buildings.

[41] T. Teixeira, G. Dublon, A. Savvides, A survey of human-sensing: Methods for detecting presence, count, location, track, and identity, ACM Computing Surveys 5 (1) (2010) 59–69.

[42] V. L. Erickson, S. Achleitner, A. E. Cerpa, Poem: Power-efficient occupancy-based energy management system, in: Information Processing in Sensor Networks (IPSN), 2013 ACM/IEEE International Conference on, IEEE, 2013, pp. 203–216.

[43] R. H. Dodier, G. P. Henze, D. K. Tiller, X. Guo, Building occupancy detection through sensor belief networks, Energy and buildings 38 (9) (2006) 1033–1043.

[44] R. Gade, A. Jorgensen, T. B. Moeslund, Long-term occupancy analysis using graph-based optimisation in thermal imagery, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2013, pp. 3698–3705.

[45] M. M. Trivedi, S. Y. Cheng, E. M. Childers, S. J. Krotosky, Occupant posture analysis with stereo and thermal infrared video: Algorithms and experimental evaluation, IEEE transactions on vehicular technology 53 (6) (2004) 1698–1712.

[46] M. M. Trivedi, T. Gandhi, J. McCall, Looking-in and looking-out of a vehicle: Computer-vision-based enhanced vehicle safety, IEEE Transactions on Intelligent Transportation Systems 8 (1) (2007) 108–120.

[47] E. Barreira, V. P. de Freitas, Evaluation of building materials using infrared thermography, Construction and Building Materials 21 (1) (2007) 218–224.

[48] F. Bianchi, A. L. Pisello, G. Baldinelli, F. Asdrubali, Infrared thermography assessment of thermal bridges in building envelope: Experimental validation in a test room setup, Sustainability 6 (10) (2014) 7107–7120.

[49] C. Balaras, A. Argiriou, Infrared thermography for building diagnostics, Energy and buildings 34 (2) (2002) 171–183.

[50] Y. Ham, M. Golparvar-Fard, 3d visualization of thermal resistance and condensation problems using infrared thermography for building energy diagnostics, Visualization in Engineering 2 (1) (2014) 12.

[51] J. R. Kominsky, J. Luckino, T. Martin, Passive infrared thermographya qualitative method for detecting moisture anomalies in building envelopes, Tedford & Pond.

[52] J. L. Lerma, M. Cabrelles, C. Portalés, Multitemporal thermal analysis to detect moisture on a building façade, Construction and Building Materials 25 (5) (2011) 2190–2197.

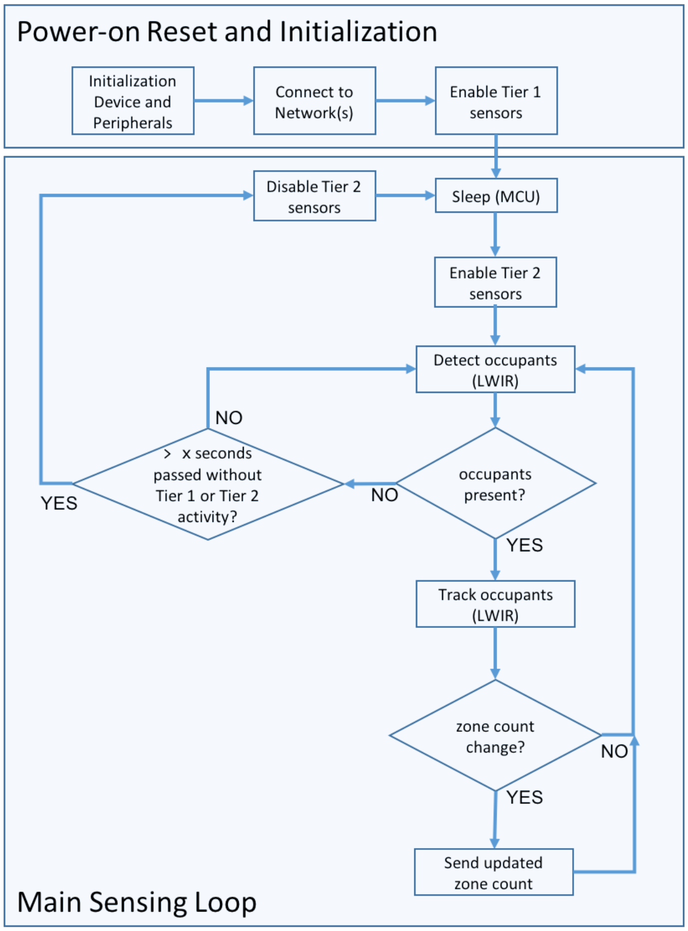

Figure 10: State diagram for sensing node operation.

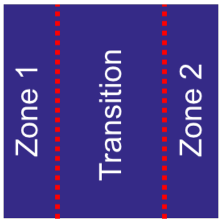

Figure 11: Logical boundaries defined within the sensor's field-of-view.

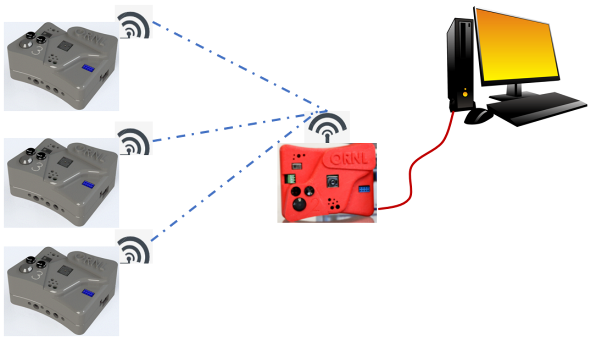

Figure 12: Communication between the nodes.

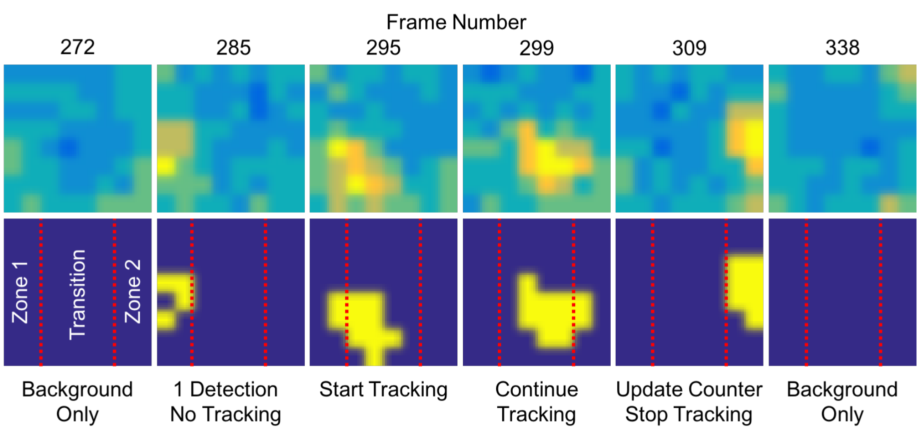

Figure 13: Raw and processed sensor output as individual modes between zones showing one person crossing from **Zone 1** to **Zone 2**.

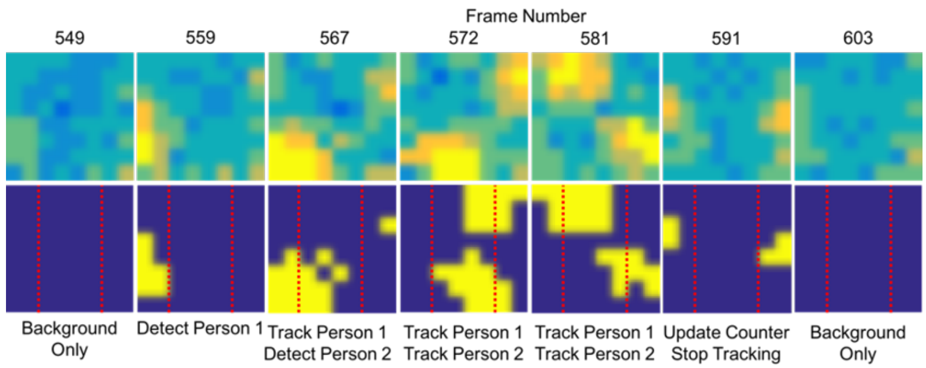

Figure 14: Raw and processed sensor output as individual modes between zones showing two persons crossing Zones (Case 1).

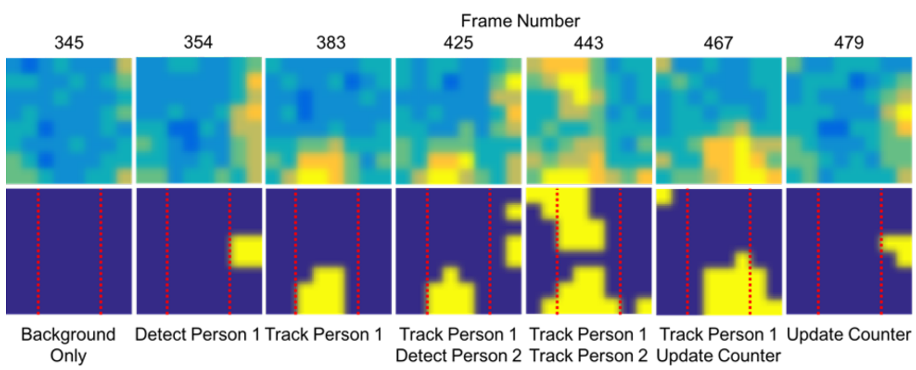

Figure 15: Raw and processed sensor output as individual modes between zones showing two persons crossing Zones (Case 2).