

A Cloud-Assisted Infrastructure for Occupancy Tracking in Smart Facilities

Dimitrios Sikeridis and Michael Devetsikiotis
Department of Electrical and Computer Engineering
The University of New Mexico
Albuquerque, NM, USA
{dsike, mdevets}@unm.edu

Ioannis Papapanagiotou
Platform Engineering
Netflix
Los Gatos, CA, USA
ipapapa@unm.edu

Abstract—The need of smarter and more interaction-friendly indoor spaces compels the combination of Internet of Things architectures with Cloud Computing approaches to produce efficient and realizable infrastructure deployments. In this work, we discuss a Cloud-assisted infrastructure model for monitoring user motion in smart buildings with reference to a real-subject realization trial. We focus on the model architecture and discuss the applicability of the generated data on Cloud-based potential application scenarios.

I. INTRODUCTION

Advanced technologies, such as Internet-of-Things (IoT) enable creative ways of enhancing productivity and life quality inside facilities. The limits are pushed beyond managing temperature, door-locks and general security, to go as far as reducing energy costs, detecting and building knowledge based on human patterns and improving the occupant-building interaction.

One of the key features such intelligent facilities should possess is the ability to efficiently track occupants' mobility, either in order to take real-time actions or use the data to calculate long-time patterns. This type of services are realised either by attempting to estimate the user's 2D coordinates in a given space, which is referred to as micro-location [1], [2], or by attempting to accurately place the user in the vicinity of certain anchor points, known as proximity sensing [3]. Especially, for indoor spaces, while a number of RF-based technologies has been proposed over the years, the unpredictability of signal propagation due to the variable physical indoor environment, makes it still an ongoing research topic [2].

Alternative methods that utilize data from body-mounted Inertial Measurement Units (IMUs) have yielded accurate results concerning microlocation [4], [5]. However, the cost of mass distribution to individuals along with the lack of a central processing and decision making node (such as a Cloud federation), prohibit their use in well-attended everyday-use environments, making them more appealing to first responder localization applications [6]. Sub-meter accuracy, scalable and low-cost installation, as well as low energy demands are important factors for large scale indoor localization services.

In this paper, we discuss a Cloud-assisted infrastructure model for location-aware smart facilities as featured in Fig. 1.




Fig. 1. Location-aware intelligent building model

The model is based on proximity-sensing and a moving transmitter-fixed scanner architecture that enables collection of user-centric data. We also provide a proof of concept-reference to our real-subject trial of this model [7] along with an occupant mobility tracking approach that leverages the collected data.

This work is organized as follows: Section II discusses related work. Section III presents the model architecture. Our proof of concept deployment and trial are presented in Section IV, while Section V discusses the basic user movement tracking method. Finally, Section VI highlights the Cloud involvement and Section VII concludes this paper.

II. RELATED WORK

Localization services based on proximity detection have been developed around a variety of architectures. Similar to our approach, Martella et al. in [8] describe a sensing infrastructure comprising of moving users equipped with transmitters, predefined broadcasting anchors and a backbone system of sniffers. Their approach is based on face-to-face and mobile-to-mobile proximity sensing representing moving users as proximity graphs [9]. Other approaches deploy less elaborate methods that utilize already installed infrastructure with the first option in such cases being the use of WiFi equipment along with smartphone deployment [10]. However, such solutions suffer from low-accuracy in distance calculations due to complex signal propagation as a WiFi beacon can be detected occasionally from buildings away.




Fig. 2. System architecture

In response to this issue, proximity sensing approaches are increasingly utilizing the Bluetooth Low Energy (BLE) technology for their services. The dominant approach in that case is utilizing BLE-beacons [11] as static short-range transmitters that mark certain areas [3], [12]–[14]. Such topology requires users equipped with scanning devices (usually smartphones) that gather the broadcast packets and use the Received Signal Strength Indicator (RSSI) values to place the user to the proximity of the anchored area. Also, these beacons are passive wireless devices without Internet connectivity or processing power, and the proximity calculations are carried out either in the user’s scanning device [13], or in most recent approaches directly to Cloud-based environments [3], [12].

III. MODEL ARCHITECTURE

Although, the aforementioned architectures can yield significant results in terms of accuracy, they heavily rely on the assumptions that users are equipped with smartphones and possess the appropriate application. Our schematic model as depicted in Fig. 1 avoids this assumption by supplying the occupants with small-factor low-consumption devices that periodically broadcast user-centric data. On the receiving end of these transmissions, we place dedicated edge devices installed throughout the facility. These scanning devices partition, either by themselves or by forming clusters, the monitored smart space into discrete cells.

At the next level of this IoT-inspired architecture, the edge devices are connected to the facility’s networking infrastructure and continuously forward the received user data to a federated Cloud level that allows centralized, but scalable coordination. That way Cloud resources can be leveraged to run proximity algorithms and provide cell to cell occupant movement detection utilizing the central supervision of the smart space. This tiered architecture will reduce the response times to the central processing Cloud deployment, resulting in a system better suited for real-time monitoring applications. Finally, the edge devices will be able to cache a significant amount of recent data before reporting to the Cloud. This capability will enhance the accuracy of the proximity estimations by utilizing all the latest measurements and will prevent data losses in the case of a temporary disconnection from the Cloud.

IV. DEPLOYMENT & PROOF OF CONCEPT TRIAL

In order to provide a proof of concept of the proposed infrastructure, a simpler version was realized at three floors of NC State University’s ECE Department as discussed in




Fig. 3. Edge device and cell locations

[7]. The deployment used 30 Raspberry Pi modules as edge devices and Gimbal BLE iBeacons Series 10 to equip the moving users. The edge devices in this case are used only as dedicated scanners that receive the beacons’ advertising packets. For every reception they report to a centralized server information regarding the identity of the beacon, the time of packet reception and its RSSI value accompanied with the scanner’s ID. The server-scanner communication is achieved via the MQTT protocol [15]. The system’s architecture is shown in Fig. 2.

The majority of edge devices was installed on the first floor of the facility and their positions were dictated by the presence of power outlets. Fig. 3 shows their positions along with their cell clusterization. Our deployment utilizes a Geofencing strategy to define localization [16]. Each edge device is considered a Point of Interest (PoI) and defines a virtual barrier around it, where our system tracks incoming and outgoing users. In our case, the geofence radius for each scanner was chosen to be seven meters to avoid overlapping areas, while for the same reason, closely installed edge devices form unified cells as shown in Fig. 3.

In order to visualize the relationship between the RSSI values and the user-scanner distances we collected and averaged 30 RSSI samples at a number of distances starting from 0 meters up to 10 meters. For this experiment the edge device was mounted on a wall at the height of 1.6 meters and the results are shown in Fig. 4. The fitted curve is based on the log-distance path loss model:

$$RSSI = -10 \gamma \log_{10}\left(\frac{d}{d_0}\right) + C \quad (1)$$

where γ represents the path loss exponent of the propagation channel, d is the user-scanner distance, d_0 is the reference distance and C represents the average value of RSSI at d_0

Finally, the installation was used to host a IRB-approved real-subject trial at the spaces of NC State University’s Centennial Campus - Engineering Building II. The trial lasted for 33 days while 46 students participated by carrying with them an iBeacon device at all times.




Fig. 4. Curve fitting for RSSI values at distances from 0 to 10 meters

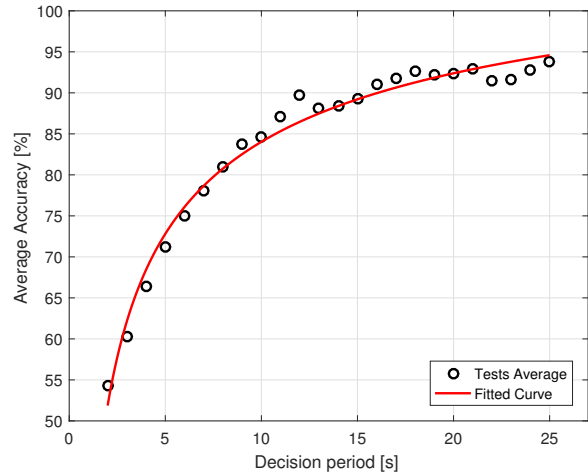


Fig. 6. Accuracy of near-edge cell occupancy - Tests average




Fig. 5. User cell transition

V. USER CELL OCCUPANCY & TRANSITIONS

Using the fixed edge devices to partition the facility spaces, proximity estimation methods can be utilised to place an individual occupant to a distinct cell. The continuous string of this user's transmitted packets provide information regarding the identity of the scanner in his vicinity and when the broadcast happened along with its RSSI. Since after a transmission burst multiple edge-devices will receive the signals, a comparing method can be used to select the cell where the user will be classified to. Using the stronger RSSI value is one of the solutions. However, we should also consider the optimal refresh period of the decision since the collection of multiple RSSI values can enable us to use RSSI smoothing filters for increased accuracy. The area covered by each scanner can be determined by simple signal propagation models that relate RSSI with distance such as Equation 1. This parameter is also tunable by considering signal strengths over a limit, reducing that way the range an advertisement packet can travel.

In order to determine the applicability of this method in locations near the cell limit we performed a series of static experiments. The classification criterion that assigns the static user to a cell is the "naive" strongest RSSI approach. The user is periodically assigned to the cell that receives the advertisement message with the strongest RSSI. The decision period was investigated in these tests with Fig. 6 showing the




Fig. 7. Generated traffic per trial day

average accuracy while the period duration ranges from 2 to 25 seconds.

By assigning the occupant to a specific and unique cell for each refresh time slot, the system will be able to track movements inside the facility as depicted in Fig. 5. It is however obvious from Fig. 6 that while higher periods yield eventually accurate results, they are not able to detect rapid user movements and therefore fast cell transitions. Therefore, we have to additionally investigate the optimal decision period to also accommodate this functionality. Over time, such data can generate useful motion patterns between cells and locations along with user-centric preferences and habits.

VI. CLOUD ASSISTANCE

Depending on the beacon advertising period, the proposed system is able to yield a significant amount of facility to Cloud traffic, especially if the facility size and user numbers are of large scale. Therefore, a Cloud federation is imperative to provide real-time services to each individual and at the same time collect the data loads for further processing. Fig. 7 shows total daily messages received by the server from all the installed edge devices during our trial.

Regarding off-line processing, the gathered data can be used not only to produce mobility patterns for each occupant,

