Estimating Dining Hall Usage Using Bluetooth Low Energy Beacons

Rachael Purta
Department of Computer Science and Engineering
University of Notre Dame
Notre Dame, IN, USA
rpurta@nd.edu

Aaron Striegel
Department of Computer Science and Engineering
University of Notre Dame
Notre Dame, IN, USA
striegel@nd.edu

UbiComp/ISWC'17 Adjunct , September 11–15, 2017, Maui, HI, USA
ACM 978-1-4503-5190-4/17/09.
https://doi.org/10.1145/3123024.3124429

Abstract
Bluetooth low energy (BLE) beacons are a relatively new technology that has recently found extensive use in various industry and marketing fields. While these beacons have begun to be used in research settings, much of the focus has been related to indoor positioning, and not on general place detection and user patterns. In this paper we use BLE beacon data to detect visits to the dining hall, and analyze these visitation patterns from the perspective of what would be useful to our institution's dining services department. For the data collection, we leverage a BLE beacon sensing-enabled institutional app to collect data from over 700 users during finals week, as a preliminary study of the capabilities of such data.

Author Keywords
User place patterns; Bluetooth low energy; beacons

ACM Classification Keywords
H.5.m [Information interfaces and presentation (e.g., HCI)]: Miscellaneous

Introduction
Since the introduction of the iBeacon protocol in 2013, Bluetooth Low Energy (BLE) beacons of all types have appeared in the market. Over time, they have increasingly become one of the essential building blocks for the future

of the Internet of Things, with a wide variety of use cases and possibilities. For example, at the Consumer Electronics Show (CES) in 2015, Radius Networks showcased the capabilities of their beacons by placing them throughout the very large exhibit area for the show's indoor navigation system. Estimote has also deployed a large amount of its beacons, using them to provide details about nearby artwork at the Guggenheim Museum.

Research with BLE beacons has, for the most part, followed the same use cases as industry, focusing on indoor positioning or proximity in some form. For example, in [4] the authors characterize the Received Signal Strength Indicator (RSSI) signal from various distance zones for the purpose of detecting how proximate a person is from an object. Positioning, specifically determining what part of the space the user is in from BLE beacons, is the focus of [3]. Other uses for beacons exist, but are less represented in research. The authors in [1] use BLE beacons to determine whether a person is traveling or stationary, although they intend to use it along with an indoor positioning system they developed. In [2], BLE beacons are used to detect users entering and exiting the office for an energy-saving algorithm. The authors here take advantage of the range of beacon detection, instead of using the more specific RSSI value from the beacon, to detect entrance and exit. In this work we use a similar technique, relying on the long range of beacon detection (5 to 50 meters, depending on how open the environment is) to detect when the user is at the dining hall.

The goal of this paper is ultimately to explore to what extent we can determine user patterns from just beacon entrance and exit information, and how useful the results may be from the operational perspective of dining services. We focused our efforts on determining the concentration of users in the dining hall at a particular hour, as well as the event patterns of when and how often a user went to the dining hall. The first analysis is useful to dining services to determine when the busiest times of day are, as well as what days are busier than others. The second analysis gives dining services more of a user perspective - meaning, it allows us to ask questions on *how* the dining hall is used, especially during an intense period like finals week.

While the results presented in this work focus on the dining hall and finals week, other locations did have beacons, though only at building entrances and exits. The extra beacons placed in the dining hall were a test in preparation for a larger deployment. For brevity, we focus here on the more dense dining hall deployment, and discuss a larger deployment in the future work section.

Methods

We collected our beacon data through a common institutional app (Notre Dame mobile, referred to as NDmobile on campus and throughout this paper) whose main purpose is to provide institutionally relevant information and services, such as athletic schedules and campus dining hours. Beacon functionality had been added to the app to collect data from various buildings with beacons, but it had not been fully explored. Beacon information from the app includes the major, minor, and time seen each time the phone entered and exited the range of a beacon. All users of the app are automatically anonymized in that they are each given a unique device identifier, but this identifier cannot be linked to any other information about them.

The NDmobile app uses Radius Networks' Proximity Kit, which provides an app-side SDK to search for and record beacons to Radius' secure servers on both iOS and Android. While we did not have access to the actual code for the app, we were informed that it used Proximity Kit and we

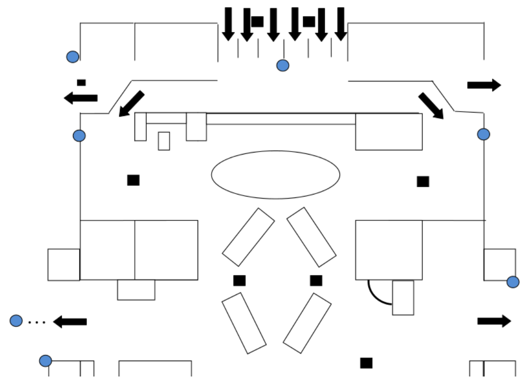

Figure 1: Dining Hall beacon placement. Circles indicate beacons, while arrows indicate traffic flow. The ellipsis indicates that the far left beacon is some space across from the arrow.

accessed the beacon data through the Proximity Kit Admin API. The Admin API does require authentication, which we obtained from the app developers at our institution. In order to make our beacons compatible with NDmobile, the Universally Unique Identifier (UUID) of the beacons was set to a value provided by the app developers. This is necessary because iOS, as well as the Proximity Kit, require predetermined UUID regions set in the code of the app. As long as we matched the UUID region of our beacons, no changes to the app were necessary.

All beacons were from Radius Networks, specifically the new X4 beacon, and all beacons had the same configuration (except for the major and minor, which together uniquely identify the beacon). Radius beacons are not required to use the Proximity Kit - all that the kit requires is the iBeacon (or Radius Networks' AltBeacon) standard and that the UUID matches the app configuration.

We installed Radius beacons at 7 strategic locations in the dining hall, focusing on entrance, exit, and transition points (see Figure 1). This is in addition to one beacon our institution had already placed for building information, not shown in the figure. The figure shows beacon placement (circles) as well as traffic flow (arrows), with the six arrows being the main entrance. Note that the focus of our experiment was the left side of the figure, as the dining area on the right side is only used for special occasions. We did, however, place two beacons on this side - one near the right entrance to the food court (near the diagonal arrow), and one near the entrance into the right dining area (rightward arrow). These two beacons allow us to cover the right side of the food court, so that users do not spend too much time out of range.

The left side beacons correspond, from top to bottom, to the transition from the hallway to the eating area (and vice versa, as this is also a common exit), the card swipe station to get into the food court, a building exit from the dining area (the beacon next to the ellipsis), and the transition from the food court to the eating area. Since our dining service collaborators did not want the beacons to be conspicuous, all beacons were placed above head height (about 7 feet or higher), which may also extend the range of beacon detection because it is more open. As our beacon data collection was a test run, our beacons were only in these locations for a total of two weeks, the first of which was finals week.

Results

We compiled beacon data from each of the days of Finals Week (May 6 to May 12, 2017) and analyzed the part of the data corresponding to the dining hall. Because of the nature of NDmobile, the users in our data can be staff, students, or even visitors (such as parents) of the university,

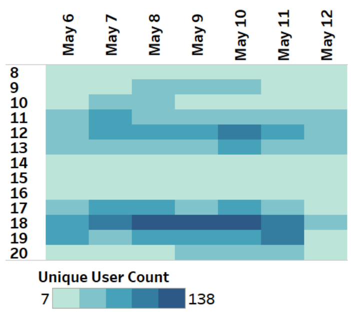

Figure 2: Number of users in dining hall during finals week, hour of day on left

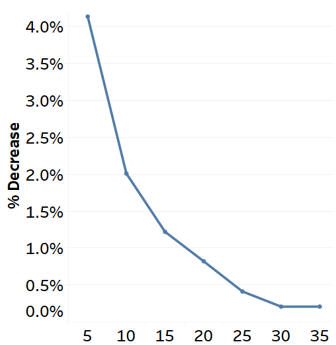

Figure 3: Percent decrease of the average number of events per user. The x axis is the gap setting, in minutes.

however users had to have their phone's Bluetooth turned on. While most of the app's users would be students, having non-students gives a more complete view of how the dining hall is used, and thus is useful to dining services from an operational perspective. In the analyses discussed below, we found 785 unique users in the dining hall during finals week.

Concentration of Users over Time

To determine the concentration of users in the dining hall, we took all starting timestamps of a dining hall beacon sighting and counted the unique number of users that appeared every hour, for every day of finals week. The unique count ensures that if a user saw more than one beacon, or saw the same beacon multiple times, they would only be counted once per hour. This gives us the user concentration heat map depicted in Figure 2, which is restricted to hours 8 through 20 (8pm) for space reasons. Note that all hours outside this range had less than 20 users, and thus were not helpful for our analysis.

In Figure 2, we can see that the highest concentration of users is around hour 18 (6pm), especially May 8 through 10 (Monday through Wednesday). The majority of exams occur on these dates, so this makes sense. We see that this pattern continues, with less users overall, on Thursday and all but disappears on Friday, where there seem to be few users in general. This is because students leave after their exams are done, and fewer exams occur on Friday. Before finals week, we see that the weekend is lower in terms of concentration than the rest of the week, Saturday being slightly lower and more dispersed than Sunday.

Figure 2 also shows the lunchtime rush occurs mostly in the 12th hour, with Wednesday the 10th having the highest number of users. This pattern persists a little more on Friday than the dinner pattern, looking more like Saturday's lunch. Overall, the lunch period at the dining hall has fewer users than the dinner one, most likely because there are many other venues for users to have meals on campus and for whatever reason, they prefer these venues over the dining hall for lunch.

Event Patterns

Since our beacon data is formatted in terms of the start and end of a beacon sighting, we wrote a script to determine the start and end of a dining hall visit event. Because of the placement of our beacons, ideally a user sees at least one beacon during the entire dining hall visit event, but in our data we noticed this is not always the case. To compensate for this, we defined an event as a period where the time gap between two instances of seeing a beacon is not longer than 30 minutes. To find this threshold, we varied the gap parameter and compared the average number of dining hall events found per user over finals week. We chose this particular measurement because the lack of seeing beacons causes the script to detect too many events per user (considering each side of the gap as a separate event), and thus the average number of events per user stabilizing means that the gap parameter is sufficient to catch most of these separated events.

The percent decrease in this average number of events per user is shown in Figure 3. As can be seen in the graph, the curve decreases almost exponentially, with the curve practically flat from 30 to 35. This means that, from the setting of 30 to 35, the percent difference has stabilized to about 0.2% decrease for every 5 minutes added on to the gap parameter. Thus we used 30 minutes as our setting for the gap parameter, as it is the lowest stable value.

Once our script determined dining hall events based on our gap threshold, we further filtered our data to not include events that were longer than 180 minutes, as this is the

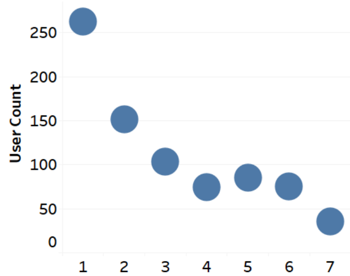

Figure 4: Number of days users went to the dining hall, with x axis as number of days

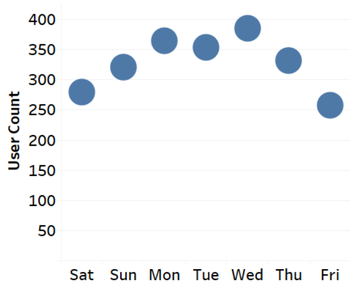

Figure 5: Days of the week users went to the dining hall, with user counts

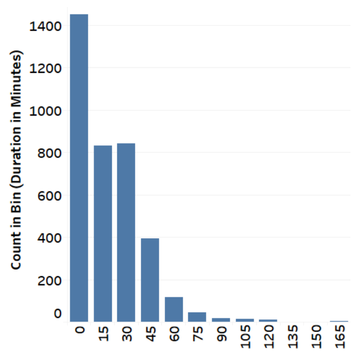

Figure 6: Distribution of the duration of all visits to the dining hall, x axis as duration in minutes

duration of lunch, the shortest of the three main meals at the dining hall (and the second most populated meal, as discussed earlier). The main reason for this is we wanted to include only users of the dining hall in our analysis, not those who had the app but were dining hall workers and not actually eating there. We do this so that our results truly reflect usage patterns of the dining hall, not working patterns, although it only filters out 2 users in this case.

The secondary reason we chose to limit to 180 minutes was because the duration distribution of all dining hall events showed no categories with more than 5 users beyond 165 minutes. We show this distribution in Figure 6. According to the figure, most of the distribution is concentrated below 45 minutes, with the 0 to 15 minute category having the most events. Since our script was optimized to catch most of the events separated from not seeing beacons, the 0 to 15 category being so large means that there are many short visits to the dining hall - possibly so students can go back to studying for exams quickly.

Our first daily usage patterns analysis examined how many days users went to the dining hall during finals week and what day of the week was most popular. The first of these is shown in Figure 4, the second in Figure 5. For Figure 4, we see that many users only went to the dining hall one day, followed by two, three, and five days, respectively. To find possible reasons why so many users appeared only one day, we looked at which days occur the most frequently among this one day category. We found that many (52) occur on Friday, the last day of exams, as well as a high number on Wednesday (43), the last day of the period that has the most exams, and Monday (40). This could mean that parents came to help their student move out and already had the app installed, or that some students are so busy during the week that they do not go to the dining hall until exams are over (with Monday users being students who finish their exams early and leave), or a combination of both. Finding out which, however, we leave to future work.

Next, we investigated which days were the most popular during finals week. The results in Figure 5 show that Wednesday (the 10th) has the most users who visited the dining hall, followed by Monday (the 8th) and Tuesday (the 9th), respectively. As discussed previously, these are the days when most exams occur, so it makes sense that a large number of students would go to the dining hall on those days. Given that Wednesday is the last of the three days and, as found in our one-day users analysis, it also has a number of users that only appeared on that day (43), it is also logical that Wednesday is the most populated day. After Wednesday, many students left for the summer, making the counts lower for Thursday and Friday.

The final question we sought to answer was 'how many times per day do users typically go to the dining hall during finals week'? We organized the data by day and the num-

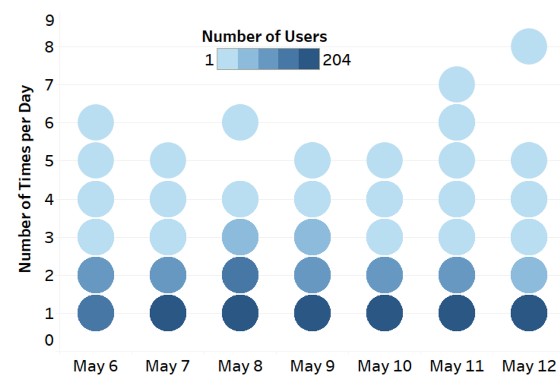

Figure 7: Times per day that users visited the dining hall.

ber of times per day, as shown in Figure 7, to find the user count for each per-day frequency. Note that the total for each day in the figure adds up to the count shown in Figure 5 for the corresponding day, as both use unique count of users per day. The first thing we observe from Figure 7 is that for each day, most users go to the dining hall only once per day. This is interesting because the only meal plans available give 14 or 21 dining hall swipes per week, allowing for 2 or 3 (or more, if saved from previous days) swipes per day. The fact that a large number of users (50% or more for each day) only use the dining hall once per day implies that there are many students who are not using all available swipes. While probably not all of these users are students, one-time users (parents and students who only visit the dining hall after finals), as discussed earlier, were at maximum 52 per day, while the one-a-day users shown here are above 150 consistently. Thursday and Friday, which have less exams, are also less populated in Figure 5, so we know that some of our users have left for the summer, as students would. This makes it likely that many of these one-a-day users are students, which is a useful trend for dining services.

Discussion and Future Work

The results discussed here are preliminary, as they only cover one week of the entire school year, and only one possible campus eating option (the dining hall). With a longer deployment and more campus dining areas covered, we could gain more insight on how students use campus dining options throughout the semester. For example, we earlier discussed finding out whether the users that only appeared once in our data were parents or students who went elsewhere for all their other meals. With more campus dining areas covered, we could answer this question.

As mentioned, we do have plans for a longer deployment. The focus of this future deployment, however, is also the dining hall, and is meant to collect more fine-grained (RSSI) information about user proximity. Because this requires a denser deployment, we plan to put over 100 beacons in the dining hall for a semester. The methods discussed in this paper still apply, fortunately, so we will be able to extend the analysis discussed here for the dining hall.

Acknowledgements

The research reported in this paper was supported by the National Heart, Lung, and Blood Institute (NHLBI) of the National Institutes of Health under award number R01HL117757. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

REFERENCES

1. Kevin Bouchard, Mahir Rafi Eusufzai, Ramin Ramezani, and Arash Naeim. 2016. Generalizable spatial feature for human positioning based on Bluetooth beacons. In *Proc. UEMCON 2016*. IEEE, 1–5.
2. Moonok Choi, Wan-Ki Park, and Ilwoo Lee. 2015. Smart Office Energy-Saving Service Using Bluetooth Low Energy Beacons and Smart Plugs. In *Proc. DSDIS 2015*. IEEE, 247–251.
3. Avgoustinos Filippoupolitis, William Olif, and George Loukas. 2016. Occupancy Detection for Building Emergency Management Using BLE Beacons. In *Proc. ISCIS 2016*. Springer, 233–240.
4. Dae-Yeob Kim, Soo-Hyung Kim, Daeseon Choi, and Seung-Hun Jin. 2015. Accurate indoor proximity zone detection based on time window and frequency with bluetooth low energy. In *Proc. MobiSPC 2015*. Elsevier, 88–95.