An Open-Space Museum As a Testbed for Popularity Monitoring in Real-World Settings

Cattani, Marco; Protonotarios, Ioannis; Martella, Claudio; van Velzen, Joost; Zuniga, Marco; Langendoen, Koen

Publication date
2017

Published in
Proceedings of the 2017 International Conference on Embedded Wireless Systems and Networks, EWSN 2017

An Open-Space Museum as a Testbed for Popularity Monitoring in Real-World Settings

Marco Cattani
Delft University of Technology
m.cattani@tudelft.nl

Ioannis Protonotarios
Delft University of Technology
i.protonotarios@tudelft.nl

Claudio Martella
Vrije Universiteit Amsterdame
claudio.martella@vu.nl

Joost van Velzen
Salland Electronics
joost.vanvelzen@salland.nl

Marco Zuniga
Delft University of Technology
m.zuniga@tudelft.nl

Koen Langendoen
Delft University of Technology
k.g.langendoen@tudelft.nl

Abstract

This paper reports our experience with crowd monitoring technologies in the challenging real-world conditions of a modern, open-space museum. We seized the opportunity to use the NEMO science center as a testbed, and studied the effectiveness of neighborhood discovery and density estimation algorithms in a network formed by visitors wearing bracelets emitting RF beacons. The diverse set of conditions (flash crowds in open spaces vs. single person booths) revealed three interesting findings: (i) state-of-the-art density estimation fails in 80% of the cases, (ii) RSS-based classifiers fail too, because their underlying assumptions do not hold in many scenarios, and (iii) neighborhood discovery can obtain exact information in an energy-efficient way, provided that static and mobile nodes are differentiated to filter out "passers by" clobbering the true popularity of an exhibit. The overall lesson from the experiment is that today's algorithms are quite far from the ideal of monitoring popularity in a privacy-preserving and energy-efficient way with minimal infrastructure across the set of heterogeneous conditions encountered in practice.

Categories and Subject Descriptors

C.2.4 [**Computer-Communication Networks**]: Distributed Systems—*Distributed Applications*

General Terms

Measurement, Human Factors, Experimentation

Keywords Crowd monitoring, Density estimation

1 Introduction

In this paper we report our experience with a system designed to monitor the popularity of different exhibits in a large science museum. Due to the size and complexity of the museum, the curators had a difficult time quantifying the interest of thousands of daily visitors on the many activities and exhibits offered by the museum.

The museum. The NEMO science center is the 5th most visited museum in The Netherlands, with half a million visitors per year. What makes this museum interesting as a case-study is not only its numbers, but rather three characteristics that are seldomly found in indoor environments. First, unlike other buildings that consist of clearly differentiated areas (rooms), the layout of NEMO can be seen as an *open 3D space* deployed over six stories with few boundaries (walls), cf. Figure 1. With more than 5,000 m^2 of exhibition space, the lack of boundaries makes it difficult to estimate the number of people participating, or interested, in each particular exhibit. Second, the challenges posed by the open-space layout are aggravated by the *large number of visitors* and their *"random" movement patterns*. NEMO is targeted mainly for kids and it does not have suggested routes. With more than 3,000 visitors per day and an average visit duration of three hours, at peak hours the museum can host more than 2,000 visitors, all roaming around with no predetermined path. Third, NEMO has *different types of exhibitions*. Some attract tens of visitors, others hundreds. Most of them are open experimentation areas where visitors spend variable amounts of time. Other events are scheduled at particular times of the day, have a rather constant duration, and can take place in either open or closed areas.

The monitoring system. The monitoring system deployed at NEMO was based on radio-frequency beacons and consisted of three main components: *bracelets* given to visitors, a network of *anchor points* placed at various points in the museum, and a network of *sniffers* to collect the data. Bracelets were constantly exchanging "discovery" beacons with similar hardware installed at each anchor point. Upon a successful encounter between a bracelet and an anchor node, a special packet describing the encounter was sent by the bracelet to the sniffer infrastructure, which committed the received information to a central database via the existing network infrastructure (Wi-Fi and Ethernet). In terms of privacy, the only requirement was for the data to be anonymous. Even though bracelets have IDs, we refrained from registering visitor details to avoid identification. For implementation details, we refer to [3].

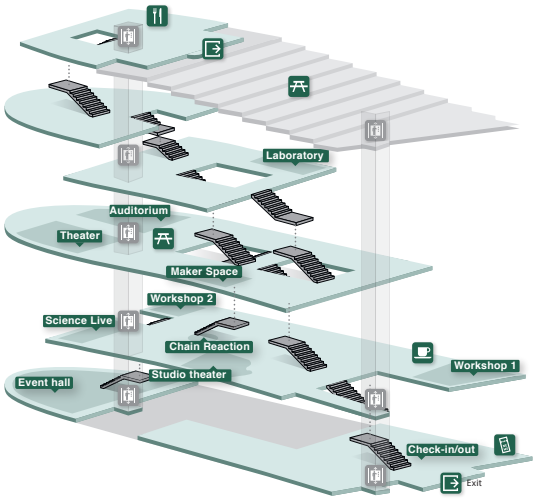

Figure 1. Floor plan of the NEMO science center.

Considering the needs and requirements of the museum's curators and the monitoring system at hand, we proposed the use of energy-efficient neighbor discovery mechanisms. Having each anchor point periodically monitor the surrounding bracelets' IDs provides sufficient information to track the popularity of an exhibit.

Being aware, however, of the growing concerns users have with ID tracking, the museum setup gave us the unique opportunity to also assess the performance of ID-less methods in a highly-unstructured scenario. Our initial hypothesis was simple: the community has developed methods to estimate node density without the need of IDs, and given that the popularity of an exhibit is determined by the number of people in its surroundings, density estimation should be an accurate proxy for popularity.

Our contribution. The evaluation shows that density estimators are ill-equipped to monitor popularity in scenarios with open spaces, high dynamics, and heterogeneous types of events. Our key findings are threefold. First, density estimators are not accurate in many cases with about 80 % of anchor points obtaining an estimation error greater than 20 %. Second, density classifiers do not work in most cases, since their underlying assumptions - lower mean and higher variance for crowded places - only apply to few application scenarios. Third, even if density estimators would be accurate, they cannot monitor popularity in scenarios having a mixture of mobile and static visitors. In these scenarios ID-less methods cannot differentiate the interested crowd (popularity) from the passersby. More work must be done to integrate the effect of crowd dynamics on ID-less methods.

2 Related Work

The analysis of human mobility is of significant importance to gain fundamental insights about people's activities, needs, and interests. Thus, many studies focus on the analysis of visitor interactions, flows, and the popularity of points of interest within buildings [8, 12, 18], fairs [14], festivals [15], and conferences [9]. For the case of monitoring and analyzing the behavior of visitors in a museum, the visualization of metrics such as popularity, attraction, holding power, and flows has been explored to support the work of museum staff [6, 13]. Many studies have also utilized sensors to measure the behavioral patterns of museum visitors. Earlier works focus on localizing visitors at coarse-grained levels (e.g., room level) through technologies like Bluetooth [17] to support multimedia guides [2, 16]. These works are conducted on traditional compartmentalized museums, focusing on classifying the visitor's experience. Our deployment provides new insights based on the challenging scenario of an open floor plan and dense crowds.

2.1 Evaluated Methods

Given the monitoring system available at the museum, we looked for methods that are amenable to RF beacons. We considered three methods: neighbor discovery, cardinality estimators, and density classifiers.

Neighbor Discovery. Neighbor discovery methods periodically monitor the set of devices in their radio range. For crowd-monitoring applications, neighbor discovery has been proved useful to provide so-called *crowd textures* abstractions to detect different crowd dynamics, like pedestrian lanes, congestions, and social groups [7], as well as to uniquely associate mobile devices with static anchor points and reconstruct the complete sequence of locations visited by a person [8].

Cardinality Estimators. Compared to neighbor discovery, cardinality estimators trade information richness for speed and efficiency. To perform their estimations, these mechanisms often model the lower communication layers, searching for features that correlate to device density. The underlying principle is based on order statistics and is simple to follow: the more neighbors beaconing, the shorter the perceived inter-beacon interval and the higher the estimated density. For example, NetDetect exploits the underlying distribution of packet transmissions [5], while Qian *et. al*. exploit the packets' arrival patterns in RFID systems for a fast estimation of tag cardinality [11]. For our deployment, we implement Estreme [4], which estimates neighborhood cardinality by measuring the average time between periodic, asynchronous beacons. We chose Estreme because of its simplicity, high reported accuracy, and minimal parameter tuning.

Density Classifiers. Density classifiers try to infer the density of a crowd by exploiting correlations with signal strength statistics [10, 14, 15, 19, 20]. These techniques require placing a set of wireless devices within the crowd, analyzing the changes of several radio features as the crowd density changes, and training a classifier. A common premise is to assume that as the density of people increases, the mean RSS decreases (due to bodies blocking radio communication) and the RSS variance increases (due to the added multipath effects). In our work we do not implement these classifiers, but focus on assessing if the RSS trends reported in the literature hold in all scenarios. (They don't.)

3 Density analysis

As stated in the introduction, resource constraints or privacy issues may prevent the use of neighbor discovery meth-

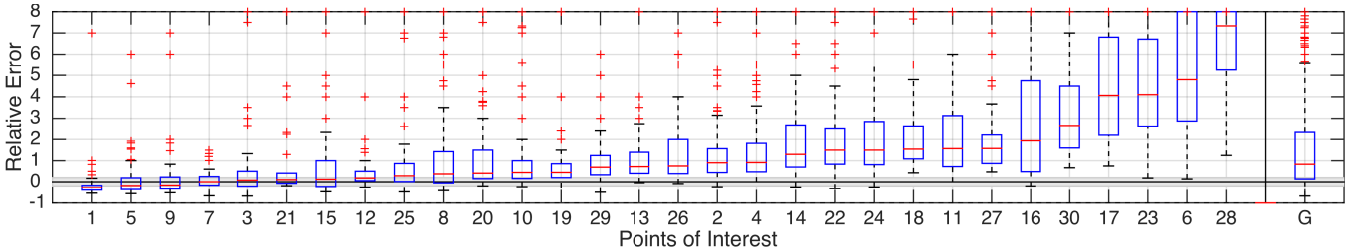

Figure 2. Density estimation error for each PoI (1…30) and for the global system (G).

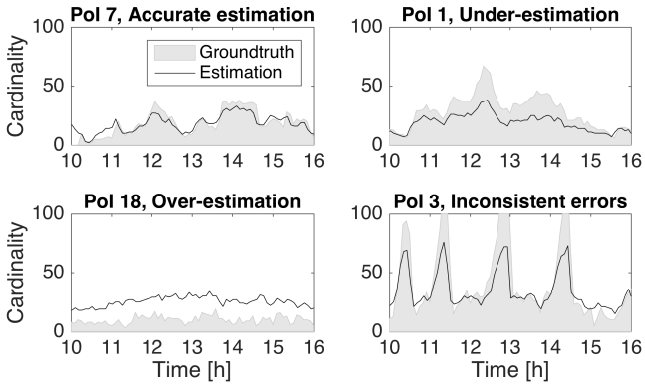

Figure 3. Density estimation errors over time.

ods in some scenarios. In these cases, density could be used as a proxy to assess the popularity of an exhibit. Thus, the first question that needs to be answered is: *"how accurate are the density estimators reported in literature in the heterogeneous setup of our museum?"*

3.1 Estimation accuracy

In our evaluation we use the information obtained with the *neighbor discovery* method as the ground truth. The error of the *cardinality estimator* is computed as $(D_{est} - D_{true})/D_{true}$, where D_{est} is the density provided by the estimator and D_{true} is the ground truth. Figure 2 shows the results per exhibit during one day of deployment, 09:00 - 17:00. The plot captures the min, max, median (red line) and the 25th and 75th percentile errors. The red plus signs depict outlier points deviating by more than two times the standard deviation from the mean. Points of Interest (PoIs) are ordered by increasing median error. The last box represents the aggregate accuracy of the global system (G). As can be seen, only a quarter of the PoIs show a median error within the range stated by the original mechanism (<20%, grey horizontal bar). Except for three PoIs, who underestimate their density, the remaining ones are all over-estimating their density, with median errors ranging from 20% to 700%. The resulting global error (G) has a median of 83%.

Why is the estimator so inaccurate? The problem lies in the fact that previous work was evaluated in a *standard building*, with *clearly demarcated areas* (rooms) and with *homogeneous densities*. In our deployment, these conditions only occur for a small fraction of PoIs, as discussed next.

3.2 Coverage versus Overlap

To better understand the causes of inaccuracy of the density estimator, we analyzed the individual estimations at each Point of Interest (PoI). Figure 3 depicts a selected representative subset of PoIs. We distinguish four common patterns.

Accurate Estimation. PoI 7 (the museum's "theater") represents an enclosed room where visitors join a workshop that runs continuously (unscheduled event). In this scenario, the conditions are similar to the ones in the original work. The density estimator performs accurately throughout the day, with errors within expectation ($< 20\%$).

Consistent Under-estimation. PoI 1 (the "event hall") represents a large open-space area where different workshops are given. Compared to the previous PoI, this is a more isolated area that is connected to the rest of the museum via a set of long stairs. This large area is covered by only one anchor point, which is unable to achieve a stable communication with some bracelets in the area (*false negatives*). Due to this poor coverage, density was consistently under-estimated, especially when the crowd increased, such as at 12:30.

Consistent Over-estimation. Due to the open 3D layout of NEMO (no barriers), the high number of people (bracelets) and dynamics (movement), we often faced the problem of overlapping coverage regions among PoIs. An example is PoI 18, an exhibit that is centrally located in the museum. Its coverage overlap leads to many false positives (duplicated IDs) and, in turn, to gross over-estimations. Under normal circumstances i.e., buildings with rooms and relatively low and homogeneous densities, false positives are less common and usually filtered out by walls.

A popular technique to discard false positives is to implement a virtual filter (virtual wall) by making estimators to discard signals below a given received signal strength (RSS). Unfortunately this filtering method did not work in our setup. To understand the reasons why, we varied the RSS threshold of the neighborhood discovery mechanism and analyzed the resulting neighborhood cardinalities (cf. Figure 4). For each moment in time, *coverage* was computed by counting the number of unique IDs in the whole system, while *overlap* was computed by summing the number of unique IDs at each PoI. As we can see in Figure 4, the more aggressive the RSS thresholding is, the smaller the overlap problem becomes. Unfortunately, at the same time the coverage problem aggravates (which leads to under-estimations), making it impossible to find an optimal thresholding value that balances coverage with overlap. The common solution to this problem is to tailor the coverage of each anchor point to the specific area of interest of the exhibit, as in [8]. Unfortunately, this option is too cumbersome to deploy in many museums, where exhibits often change and require constant adjustments.

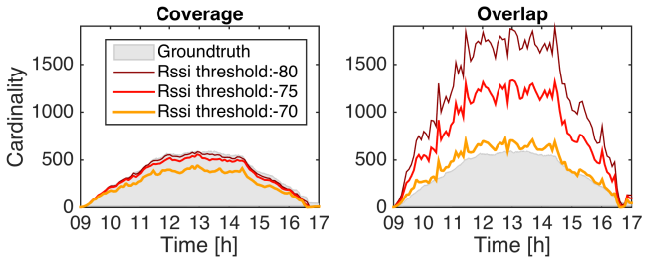

Figure 4. Coverage and overlap for different RSS values.

Inconsistent Estimation Errors. While for the previous two types of estimation errors it is possible to solve the coverage/overlap problem with a careful and time consuming deployment, for exhibits with large coverage areas and highly heterogeneous densities even a perfect positioning and filtering system would not reduce the inaccuracies.

This is the case for PoI 3 (the museum's "chain reaction" show), which is the most famous exhibit in the museum, and spans from the first to the third floor of the building. This show is repeated 5 times a day, lasts 15 minutes, and attracts a large fraction of the visitors around the internal balcony of the museum. As shown in Figure 3, the estimation is quite accurate during the time where there is no show (low homogeneous density composed by people passing by), but during the show the density is heavily under-estimated.

This estimation error is not due to a coverage problem, but is an artifact of how estimators cope with sudden changes in density. The estimation process is sped up by averaging local estimates with neighboring ones, effectively increasing the number of samples, in turn boosting convergence. An adverse effect is that this approach smooths the spatial density peaks, leading to under-estimation in denser areas (PoI 3) and over-estimation in sparser ones.

3.3 Effects of Crowd Dynamics

We now analyze the effect of crowds on the propagation of radio signals. This analysis highlights the limitations of density classifiers. Our deployment has 11 static bracelets. We used these bracelets, together with anchor nodes, to monitor the changes in radio signals throughout the day.

Figure 5 depicts the effect that density has on radio coverage and signal strength. Every 300 s, each *static* point (bracelet and anchor) calculated (i) the average density in its surroundings, (ii) the average RSS from other *static* points and (iii) the average "radio range". The average radio range was obtained as follows: upon receiving a packet from a static point, the receiver calculated its distance to the sender (positions of static points are known), and averaged this distance with all other distances observed during the last 300 s. The boxplots represent the averages during a day.

Effects on Range. The results in Figure 5 (left) indicate that an increase in crowd density reduces the maximum range of devices, while not significantly affecting the median. This is not too surprising, since people are known to cause signal attenuation due to body effects, which filters out weaker radio signals sent by distant nodes. The result does indicate, however, that existing cardinality estimators are ill-equipped to calculate densities (number of devices over a given area),

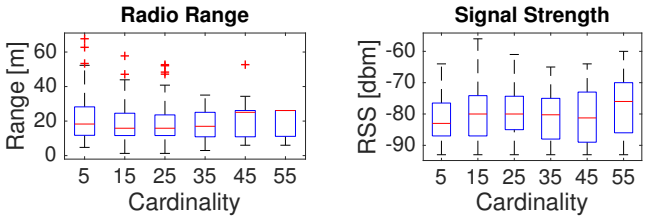

Figure 5. Effect of the crowd density on static nodes.

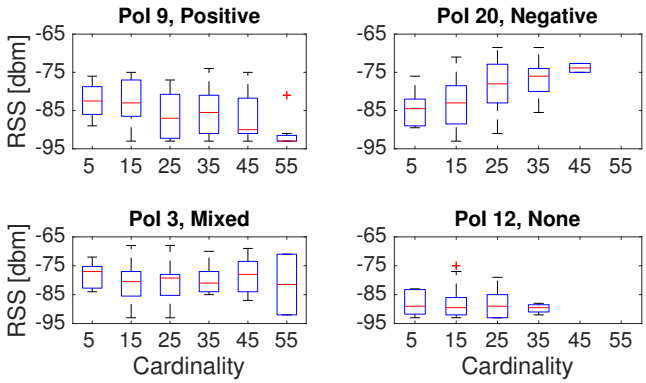

Figure 6. Effect of the crowd density on RSS.

since the radio range is oddly influenced by the crowd.

Effects on Signal. According to many crowd density classifiers, when the crowd density increases, signal strength characteristics should get worse. Similar to what is hypothesized for radio range, the assumption is that an increasing number of people should reduce the average RSS (due to body effects) but increase its variance (due to multi-path effects).

Surprisingly, this relation does not hold for the RSS values observed in our deployment. Figure 5 (right) shows no clear correlation between crowd density and key features in the RSS (median and variance). To better understand why our observations diverge from previous works, we analyzed the relation between density and signal strength for each PoI. Figure 6 depicts a representative subset. Note that each PoI in this subset showcases a different relation.

Negative correlation. PoI 9 (an open-space exhibit) has a rather constant and homogeneous density throughout the day, cf. Figure 2. In this scenario, the relation between density and RSS follows the assumption made in prior studies, with a decreasing median and a slightly increasing variance (except for density 55, for which we do not have enough data points to accurately compute the variance).

Positive correlation. Surprisingly, even though PoI 20 represents a scenario similar to PoI 9, its RSS values show the opposite relation with density. We do not know the exact reasons for this difference, but we hypothesize that it is caused by the different levels of interference perceived by the two PoIs. Figure 2 shows indeed that PoI 20 is affected by more false positives than PoI 9. This implies that PoI 20 is exposed to more overlapping areas and long links, which increase interference. According to [1], this increased interference could lead to RSS values that are up to 8 dBm higher than normal, an offset that is almost as big as the range observed in Figure 6.

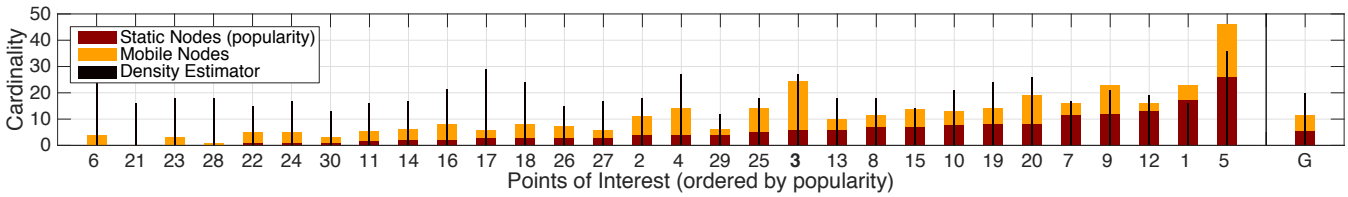

Figure 7. Breakdown of static and mobile nodes (neighborhood discovery) and the density estimates.

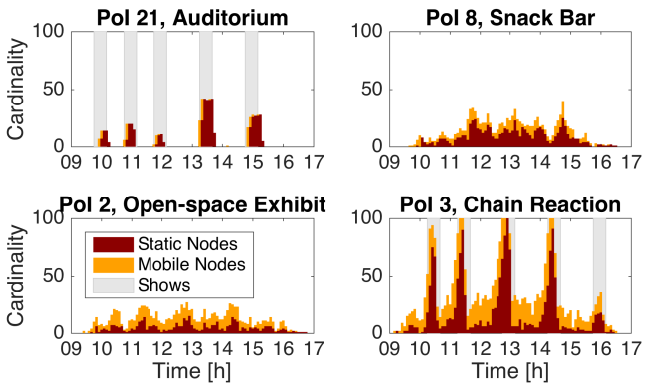

Figure 8. Density and popularity over time.

Inconsistent correlation. PoI 3 (the open-space "chain reaction" show) presents inconsistent behavior, with median RSS values that both increase and decrease with density. This bimodal behavior is probably caused by the peculiar characteristics of this PoI, which periodically morphs from a point of transit to a crowded point of attraction (and vice versa).

No correlation. Finally, PoI 12 (a laboratory setting) features a limited number of participants, all seated and well spaced from each other. In this scenario the crowd is so sparse and small that it does not affect the perceived signal strength.

As shown the relation between signal and density drastically changes depending on the scenario. Existing density classifiers work only when assuming a uniform distribution that does not change drastically over time. We argue that for such classifiers to work in heterogeneous scenarios, such as our museum, PoIs would need to undergo cumbersome training to account for the peculiar radio features of each area.

4 Popularity Analysis

As was shown in the previous section, density estimation and classification techniques struggle with the heterogeneous set of conditions found in the museum, producing inaccurate results in many cases. What is worse, in this section we will show that density is a poor proxy for the popularity of an exhibit. The reason is that density does not discriminate visitors engaged with the exhibit's activity from "passers by" that happen to be in the vicinity for a short while. This is especially a concern in our test-case museum as the lack of fixed routes (e.g., corridors) in combination with the open-space layout makes it hard to physically separate visitors and passersby, which rules out the use of RSS thresholding to limit the area of interest surrounding an exhibit. The concept of an area of interest is also dynamic, as at busy times people queue up in unpredictable ways.

The definition of popularity that we will be using throughout the remainder of the paper is the number of visitors that stays in the vicinity of an anchor point for some time window W. More formally, we distinguish Static (S) and Mobile (M) nodes:

$$\mathrm{S}(i,t) = \bigcap_{w=0}^{W} n_i^{t-w}, \quad \mathrm{M}(i,t) = n_i^t - S(i,t),$$

where n_i^t is the neighborhood discovered by PoI i at time t. The popularity of exhibit i at time t is then simply given by $|S(i,t)|$. We experimented with lengths of window W from 10 to 300 seconds. As longer windows provide a smoother display of changes over time, we will present the results from the largest window size (W = 300 s) in the sequel of the paper. With shorter windows, similar results were obtained.

The importance of properly discriminating static and mobile nodes can be seen in Figure 7, which plots the median number of static (red bars) and mobile (yellow bars) nodes observed with our neighborhood discovery data during a day. The higher the red bar, the more popular the PoI is. Note that the fraction of mobile nodes differs considerably between PoIs, and is uncorrelated with popularity. Overall, about half the nodes are mobile, and should be disregarded. For example, looking at the aggregated cardinality, PoI 3 would be the second most popular exhibit, while in fact it is only the 11th most popular one (PoI 1 is the second most popular). For reference, Figure 7 also includes the cardinalities produced by the density estimator (black lines), showing that using them would result in an even-more incorrect ordering of the popularity of the PoIs.

While Figure 7 presents the median cardinalities across a whole day, it is also instructive to look at how the fraction of mobile nodes changes over time as interesting patterns can be observed. Figure 8 plots the static/mobile breakdown for four exemplary scenarios.

Points of attraction. PoI 21 (the auditorium) is enclosed by walls and features a periodic shows that attracts tens of visitor for a short, fixed period of time (grey bars in Figure 8). In these conditions the flow dynamics are minimal as people enter the room shortly before a show starts; note the minimal yellow shading at the left of each show. Consequently, density estimations maps almost perfectly to popularity for the auditorium. PoI 8 (a small bar outside the auditorium) monitors visitors having a short snack (intended) as well as the ones passing by (unintended) clobbering the real popularity.

Point of passage. PoI 2, on the other hand, represents an exhibit located in an open space, close to a point of passage i.e., an area that is similar to a corridor but does not have walls. In this case most of its perceived density it is due to people passing by. To understand the popularity of PoIs like this, it is therefore essential to distinguish between static and

dynamic neighboring bracelets, especially since the fraction of mobile nodes varies across time.

Temporal hybrids. Finally, PoI 3 (the "chain reaction") is normally a popular point of passage (amidst several staircases, see Figure 1), but periodically becomes a point of attraction for a short duration (show times are highlighted with a grey background). The error, i.e., the number of moving nodes taken into account, in this case drastically fluctuates, with rather accurate estimations during the periodic shows (hardly any yellow on top) and bad estimates during the remaining periods (lots of yellow).

5 Fun Facts

Although experimentation with a thousand devices comes with a lot of hard work and frustration, the museum experience also revealed some fun facts that kept us going.

- Independently of the battery capacity, the lifetime of a bracelet is just two days. We argue that this is due to the fact that kids like to use bracelets as swords.
- Having a beaconing mechanism is useful when recovering bracelets abandoned at the most unusual places i.e., plant pots and urinals.
- Ethernet is not as reliable as WiFi; wall sockets can easily be powered down by technicians heading home at 5PM.

6 Conclusions and Future Directions

This paper reported our experience in monitoring the popularity of exhibits in a modern, open-space science museum. We equipped visitors with bracelets emitting RF beacons, allowing us to study crowd-monitoring mechanisms such as neighbor discovery algorithms, cardinality estimators, and density classifiers in a diverse set of real-world settings.

An in-depth analysis of the collected data (proximity detections and density estimations) provided a better understanding of the limitations of the existing mechanisms. In particular, the use of ID-based neighbor discovery allowed us to univocally associate each bracelet to the point of interest, and distinguish static from mobile nodes. The latter proved key to understanding the relation between density and popularity. Cardinality estimators and density classifiers cannot provide such differentiation, resulting in estimation errors far higher than originally reported (more than 20% error in 80% of the cases). We argue that these mechanism could be easily improved by letting bracelets sense their mobility, e.g. with accelerometers, or by letting anchor points exchange information and distributedly disambiguate each bracelet. In the case of density classifiers, similar improvements could be achieved by developing more complex models capable of differentiating their estimations based on the application scenario (open-space, rooms, ...). An important line of future research is thus to devise lightweight and privacy-preserving crowd-monitoring algorithms that can univocally associate each person to a single anchor point while differentiating between static and mobile visitor. Only in this way one will able to provide a precise estimation of popularity.

7 Acknowledgments

We would like to thank the anonymous reviewers and our shepherd Yuan He for providing us detailed feedback on draft versions of this paper. This work was supported by the Dutch national program COMMIT/ (project P09EWiDS) and NEMO's Science Live research program.

8 References

[1] A. Behboodi, N. Wirstrom, F. Lemic, T. Voigt, and A. Wolisz. Interference effect on localization solutions: Signal feature perspective. *IEEE Vehicular Technology Conference*, 2015.

[2] E. Bruns, B. Brombach, T. Zeidler, and O. Bimber. Enabling mobile phones to support large-scale museum guidance. *IEEE Multimedia*, (2):16–25, 2007.

[3] M. Cattani. *Opportunistic Communication in Extreme Wireless Sensor Networks*. PhD thesis, Delft University of Technology, 2016.

[4] M. Cattani, M. Zuniga, A. Loukas, and K. G. Langendoen. Lightweight Cardinality Estimation in Dynamic Wireless Networks. In *13th ACM/IEEE Int. Conf. on Information Processing in Sensor Networks (IPSN)*, pages 179–189, 2014.

[5] V. Iyer, A. Pruteanu, and S. Dulman. NetDetect: Neighborhood Discovery in Wireless Networks Using Adaptive Beacons. In *5th Int. Conf. on Self-Adaptive and Self-Organizing Systems*, pages 31–40, oct 2011.

[6] J. Lanir, P. Bak, and T. Kuflik. Visualizing proximity-based spatiotemporal behavior of museum visitors using tangram diagrams. In *Computer Graphics Forum*, volume 33, pages 261–270. Wiley Online Library, 2014.

[7] C. Martella, M. Dobson, A. V. Halteren, and M. van Steen. From proximity sensing to spatio-temporal social graphs. In *IEEE Int. Conf. on Pervasive Computing and Communications (PerCom)*, pages 78–87, 2014.

[8] C. Martella, A. Miraglia, M. Cattani, and M. V. Steen. Leveraging Proximity Sensing to Mine the Behavior of Museum Visitors. In *IEEE Int.Conf. on Pervasive Computing and Communications (PerCom)*, 2016.

[9] C. Martella, M. van Steen, A. Halteren, C. Conrado, and J. Li. Crowd textures as proximity graphs. *IEEE Communications Magazine*, 52(1):114–121, 2014.

[10] M. Nakatsuka, H. Iwatani, and J. Katto. A Study on Passive Crowd Density Estimation using Wireless Sensors. *4th Intl. Conf. on Mobile Computing and Ubiquitous Networking (ICMU)*, (2):1–6, 2008.

[11] C. Qian, H. Ngan, Y. Liu, and L. Ni. Cardinality Estimation for Large-Scale RFID Systems. *IEEE transactions on parallel and distributed systems*, 22(9):1441–1454, 2011.

[12] A. J. Ruiz-Ruiz, H. Blunck, T. S. Prentow, A. Stisen, and M. B. Kjaergaard. Analysis methods for extracting knowledge from large-scale wifi monitoring to inform building facility planning. In *IEEE Int. Conf. on Pervasive Computing and Communications (PerCom)*, pages 130–138, 2014.

[13] R. Strohmaier, G. Sprung, A. Nischelwitzer, and S. Schadenbauer. Using visitor-flow visualization to improve visitor experience in museums and exhibitions. In *Museums and the Web (MW2015)*, 2015.

[14] J. Weppner and P. Lukowicz. Collaborative crowd density estimation with mobile phones. In *3rd Int. Workshop on Sensing Applications on Mobile Phones (PhoneSens)*, 2011.

[15] J. Weppner and P. Lukowicz. Bluetooth based collaborative crowd density estimation with mobile phones. In *IEEE Conf. on Pervasive Computing and Communications (PerCom)*, pages 193–200, 2013.

[16] G. Wilson. Multimedia tour programme at tate modern. In *Museums and the Web*, volume 3, 2004.

[17] Y. Yoshimura, F. Girardin, J. P. Carrascal, C. Ratti, and J. Blat. New tools for studying visitor behaviours in museums: a case study at the louvre. In *Int. Conf. on Information and Communication Technologies in Tourism (ENTER 2012)*, pages 391–402, 2012.

[18] C.-W. You, C.-C. Wei, Y.-L. Chen, H.-H. Chu, and M.-S. Chen. Using mobile phones to monitor shopping time at physical stores. *IEEE Pervasive Computing*, 10(2):37–43, 2011.

[19] Y. Yuan, C. Qiu, W. Xi, and J. Zhao. Crowd Density Estimation Using Wireless Sensor Networks. In *Int. Conf. on Mobile Ad-hoc and Sensor Networks*, pages 138–145, 2011.

[20] Y. Yuan, J. Zhao, C. Qiu, and W. Xi. Estimating Crowd Density in an RF-Based Dynamic Environment. *IEEE Sensors Journal*, 13(10):3837–3845, 2013.