Locating Emergencies in a Campus Using Wi-Fi Access Point Association Data

Asma Ahmad Farhan
Computer Science &
Engineering Department
University of Connecticut
Storrs, CT 06269 USA

Athanasios Bamis
Computer Science &
Engineering Department
University of Connecticut
Storrs, CT 06269 USA

Bing Wang
Computer Science &
Engineering Department
University of Connecticut
Storrs, CT 06269 USA

UbiComp'13 Adjunct, September 8–12, 2013, Zurich, Switzerland.
Copyright © 2013 ACM 978-1-4503-2215-7/13/09...$15.00.

http://dx.doi.org/10.1145/2494091.2499579

Abstract

Despite much progress in emergency management, effective techniques for real-time tracking of emergency events are still lacking. We envision a promising direction to achieve real-time emergency tracking is through widely adopted smartphones. In this paper, we explore the first step in achieving this goal, namely, locating emergency in real time using smartphones. Our main contribution is a novel approach that locates emergencies by analyzing AP (access point) association events collected from a campus Wi-Fi network. It is motivated by the observation that human behavior and mobility pattern are significantly altered in the face of emergency, which is reflected in how their smartphones associate with the APs in the network. More specifically, our approach locates emergency by discovering APs with abnormal association patterns using Extreme Value Theory (EVT). Preliminary evaluation using real data collected from a university campus network demonstrates the effectiveness of our approach.

Author Keywords

Emergency management, Wi-Fi networks, Emergency tracking, Smartphones, Extreme Value Theory (EVT)

ACM Classification Keywords

C.2.3 [Network Operations]: Network monitoring

Introduction

Campuses have traditionally been in the epicenter of safety incidents. Perhaps the most publicized events involve attacks aiming at mass murder. Recent occasions include the Virginia Tech and the Columbine High School massacre, the Sandy Hook Elementary shooting and many more. Besides these major scale incidents, other more frequent emergency situations include among others: fires, bomb threats, riots, accidents, and natural disasters. All these public safety risks have sparked significant interest in methodologies and policies for emergency management [24, 25], as well as a plethora of software platforms and tools [21, 16]. These policies and tools, however, focus primarily on either preparing for emergencies or responding after an incident has occurred. In most cases information about the evolving situation is collected through phone calls and surveillance equipment available near the location of the event. Even though phone calls usually occur within minutes from the occurrence of an incident, they typically offer only fragmented information of the ongoing situation, leading to confusion and delays in emergency response. Similarly, despite significant progress in surveillance equipment and video analytics (see recent survey [14] and references thereof), the high costs and the difficulty of human operators to monitor surveillance equipments in real time often limits their use as a tool for post-event analysis [10].

We envision a promising direction in tracking emergencies in real time is through widely adopted smartphones. This is because, smartphones, being physically small and constantly carried by their owners, are effective “human sensors.” When connected to large-scale Wi-Fi networks (e.g., university, corporate, government or research campus networks), through the access points (APs) that they are associated with, smartphones, and hence their owners, can be automatically located in real time. In the face of emergency, real-time location of people is immensely valuable for rescue efforts, and also provides valuable real-time information on where the event originates and how the event evolves over time.

In this paper, we explore locating emergency through smartphones in real time, the first step to achieve our ultimate goal of real-time tracking of emergencies. Localizing and tracking emergencies in real-time is particularly important, because even though people report them immediately through 911 calls and social media (e.g., twitter), the information is often fractured or inaccurate, particularly under the chaotic conditions in the face of an emergency. Our intuition is that in emergency situations the behavior and mobility patterns of people present in the scene will be significantly altered. By discovering APs that have “abnormal” association patterns, the location of the emergency can be determined. Discovering “abnormal” association patterns, or “outliers”, is, however, a particularly challenging task because of the highly stochastic nature of human behavior. For example, small delays in the duration of a class, relocations of meetings or a workshop taking place in a building could all potentially generate a large number of outliers. Moreover, seasonal patterns that are prevalent in most campuses dictate that the developed models should be able to rapidly adapt to permanent or transient changes. To make things worse, the limited number and diversity of emergencies that typically occur in a campus imply that any model will have to be trained using only positive samples (i.e., instances of normal occupancy patterns). To deal with the above challenges, we propose a method for detecting “abnormal” association patterns using Extreme Value Theory (EVT) [22].

Our contribution in this paper is a method for detecting emergencies by analyzing AP association events collected by a Wi-Fi network. The proposed method uses network information and traffic patterns to identify smartphones. Subsequently, the AP association events generated by smartphones are used to extract features that describe the behavior of people at different areas of the campus (e.g., buildings) for different time intervals of the day (e.g., between 9am-9:15am). The extracted samples (feature vectors) are used to learn the parameters of a distribution that describes the "normal" state of the campus. Using the principles of EVT, an "extreme" distribution is extracted from the "normal" data, and is used to set a "threshold" above which an event can be classified as "emergency". To evaluate our approach we use data collected from the Wi-Fi network of our university over a period of 5 months and emergency logs from the fire department. Preliminary results show low false alarms and high detection ratios, demonstrating that our approach is promising.

Related Work

Because of a number of recent large-scale incidents, information and communication technology has gained an important role in emergency response and crisis management [24]. Perhaps the most commonly used technology today concerns surveillance equipment. The majority of campuses are currently equipped with camera networks that provide live feeds from different areas to human operators. In spite of the huge leaps that computer vision has made in automated video analytics, scene understanding and particularly crowd behavior analysis, most existing commercial systems still require significant effort from their operators. The major drawback of existing computer vision based approaches is that they are designed to operate under predefined circumstances, and lack flexibility. When it comes to emergency detection, the problem is further intensified, because no two emergencies will appear the same. This makes designing rules or supervised machine learning approaches for detecting anomalies particularly challenging. Due to these reasons, most of the computer vision based techniques are used for post-facto analysis [11]. Additionally, due to the associated installation, maintenance and operation costs, as well as due to their intrusive nature, camera networks are not suitable for use in all areas, and especially inside buildings.

Thanks to the growing number of smartphone users, several applications have emerged aiming to provide campus occupants with preparedness tools and/or real-time alerts [21, 16]. Most of these tools, however, are used for communicating information to the users rather than collecting information about the emerging situation. In this paper, we propose to complement existing surveillance systems with implicit information about the user' locations collected from smartphones. Most campuses are already covered by ubiquitous Wi-Fi networks, that cover most indoor areas, and often most outside areas of a campus. Access Point (AP) association events are already collected by the administrators of these networks for use in various network management issues [5]. Our goal is to use this already available information to detect emergencies in real-time with the goal to improve the effectiveness of first responders.

To detect emergencies, our approach seeks to detect outliers in features extracted from the network's AP association events (e.g., number of people exiting a building). Outliers detection in streams of sensor events is a relatively common problem in data mining and specifically when using sensor and/or network data. For example, a virus detection and alert system based on statistical analysis of network traffic is proposed in [7]. Similarly, entropy-based anomaly detection metrics are introduced by [17, 19] and graph-based data mining methods are introduced by [23, 4]. Analyzing human behavior, however, is significantly more challenging because of the highly stochastic nature of humans and the different seasonal patterns that they tend to follow. In this paper, we focus on unsupervised methods to detect outliers. Observing that emergencies are rare events, we believe extreme value theory (EVT) is a natural framework for the problem we are trying to solve. Therefore, we apply EVT to automatically detect emergencies in this paper; exploring the usage of other outliers detection techniques is left as future work.

Even though a lot of work exists in analyzing human behavior using sensor data, especially for wellness applications [2], when it comes to Wi-Fi network meta-data, there is only a limited amount of relevant work. In [18] Moghaddam et al. used information theoretical co-clustering to extract the relationship between users and their website access patterns. Through analysis, it was shown that the website access pattern varies with location and device. In a different approach Hsu et al. [13] used singular value decomposition and hierarchical clustering of the users' web site access pattern to extract behavioral models and find similarities among them. Finally, the work by Kumar et al. in [15] attempts to identify the gender of a Wi-Fi networks' users by analyzing their behaviors. Unlike our goals, the majority of this work seeks to identify common patterns across users or extract high-level aggregate information. In our case, however, our goal is to detect emergency events. The fact that these events are extremely rare makes the use of any supervised learning approach impossible. Furthermore, most common statistical models are inadequate for identifying outliers in a campus environment due to the seasonal patterns (e.g., daily, weekly, semester) and various random events that often take place. More specifically, the inherent randomness in the Wi-Fi users' behavior means that approximating the underlying probabilistic distributions is a particularly hard problem.

Problem Description

Our goal is to locate emergencies on a campus (e.g., university, corporate, government or research campus) in real time. Specifically, let $\mathcal{A}$ denote the set of areas that are of interest, where an area $A \in \mathcal{A}$ can be a building, or a wing/floor of the building. Our goal is to determine whether there is an emergency in area A in real time. On the high level, our approach is to learn the typical occupant behavior at area $A \in \mathcal{A}$, and use it to detect anomalies that could indicate an emergency in this area. The intuition is that in cases of emergencies like fire and bomb threats, the typical behavior of the occupants in A will be significantly altered. For instance, people may move from one wing of a building to another wing, or try to evacuate the building. These movement patterns will be reflected by how their smartphones connect to the campus Wi-Fi network, specifically, how they are associated with the APs. To this end, our goal is to discover anomalies in the AP association events of an area that could potentially indicate the occurrence of an emergency. In the following, we first describe our data collection strategy and then provide an overview of our overall approach for detecting anomalies.

Data Collection

Consider a campus Wi-Fi network. A data collection server periodically polls AP association events from all the APs using a standard network protocol (e.g., SNMP).

Each association event is denoted as $e_i = (a_i, m_i, t_i, d_i)$, where a_i is the AP to which the device with MAC address m_i is associated with, t_i is the time that the association is established, and d_i is the duration of the event. Since different types of wireless devices may be connected to the Wi-Fi network and we are interested in smartphones (because they are more correlated with human locations than other types of devices such as laptops), we need to identify the AP association events that are generated by smartphones. Determining whether an association event corresponds to a smartphone or not is achieved by checking whether the device MAC address, m_i, is in a database that stores smartphone MAC addresses. The database is obtained by analyzing DHCP logs collected from DHCP servers (since smartphones and non-smartphones have different fingerprints in their DHCP messages) [20], or by analyzing various keywords in HTTP requests [1, 12, 6], which is executed periodically to keep the database up to date. After identifying AP associate events from smartphones, we obtain high-level features (e.g., number of exits) for each AP using a sliding time window (see more details later in this section). *To increase user privacy, only aggregate information for the features (e.g., total number of exits in a window) is stored by our system.* While the data thus collected is opportunistic in nature, in that it can only track people when their smartphones are connected to the Wi-Fi network, we believe it can still provide significant insights about their collective behavior, and is hence useful for locating emergencies. As we will see, this is confirmed by our evaluation results using real network data.

High-level Approach

As mentioned earlier, we determine whether emergency occurs in an area by detecting significant changes in its occupants' behavior. In the face of an emergency, we expect at a minimum changes in the number of devices entering or exiting an area due to people moving away from the area of a threat, or conversely moving into an area to find shelter or provide help. These changes are reflected by two features of the AP association data, namely, the number of association and disassociation (or exit) events at an AP. To detect emergency, we need to learn the "normal" values of these features at different times of the day, and use this knowledge to detect outliers. The main challenge is that different parts of the campus can have drastically different occupancy patterns depending on the time of the day, the day of the week, and even the season of the year (e.g., fall, summer). To deal with this issue, we split days into a set $\mathcal{W} = \{w_0, \ldots, w_n\}$ of fixed duration windows (e.g., 15-minute windows), where w_i is an interval (e.g., $w = [9\text{am}, 9:15\text{am}]$), and obtain the features during each such time window. Then for each season, weekday and time window, we attempt to learn the distribution of the features in the area. In the rest of the paper, we only consider a single feature, i.e., the number of exits; considering other features is left as future work. The number of exits in area A during time window $w \in \mathcal{W}$ is denoted as $X_{A,w} = |\{e_i \mid a_i \in A \text{ and } t_i + d_i \in w\}|$, where with a slight abuse of notation, we use A to represent the set of APs in area A.

We treat the number of exits $X_{A,w}$ as a random variable. One might be tempted to model the random variable using a distribution (e.g., normal) and use a static threshold to detect outliers. For instance, one may choose the threshold to be the 99th percentile in the current distribution. The major drawback of this approach is that, as the number of training samples increases, there is a natural increase in the number of samples that would be above the threshold, leading to increased false alarms, which is unacceptable in critical applications like emergency detection.

Essentially, our inability to select a proper threshold stems from the fact that the vast majority of samples used to learn the distribution of the number of exits belong to the "normal" class. This happens because by definition emergencies are extremely rare events. The distributions used to model "normal" data focus more on the central tendency of the data, since the majority of the available samples is centered around the mean value. In our case, however, we are mainly interested in the tail parts of the distribution that can be used to detect emergencies. To study the tails of the distribution, we use Extreme Value Theory (EVT) that provides a framework to study the behavior of tails of a distribution, and has been successfully used to model rare events such as extreme weather or stock market crashes.

Given samples of a random variable, in our case $X_{A,w}$, EVT attempts to assess the probability of events that are more extreme than any previously observed sample. The key idea behind it is to use the extreme values (i.e., extracted through block maxima) of a random variable to create a new tail distribution that can more accurately describe the probabilities of the extreme events. Subsequently, an outlier can be detected by setting proper thresholds in this new extreme value distribution. A more detailed overview of EVT is provided in the following section.

Locating Emergencies Using EVT

To locate emergencies we seek to determine whether the number of exits from an area (represented by the number of AP disassociation events in the area) is an outlier or not. As described earlier, we use EVT to detect outliers. In the following, we first provide more details of EVT, and then describe how we can use it to detect outliers.

Extreme Value Theory (EVT)

EVT is a statistical framework used to model the extreme tails of a distribution [8, 22, 9]. The main advantage of EVT is that it can fit the extreme values of the data regardless of the overall original distribution. This allows to estimate the distribution of maxima without the need to learn the underlying parent distribution. EVT is particularly suitable in our case, because of the constantly changing and unpredictable occupancy patterns of most campus buildings. This is true not only due to the differences in the number of occupants observed across days (e.g., due to differences in class schedules or events) but also due to seasonal patterns that most campuses tend to follow.

In this work, we use Generalized Extreme Value (GEV) distribution to model the extreme values of the number of exits in a building. To learn the GEV distribution, the block maxima approach is used, which divides the features into contiguous blocks. Specifically, let $S_{A,w}$ be the features extracted (i.e., exits) from the association events. We split our input S into contiguous blocks of size k and for each such subset we pick the maximum value. Using block maxima approach, the maximum would be represented as:

$$x_k = \max\{x_1, x_2, x_3, \ldots, x_k\} \tag{1}$$

where k is the size of a block.

The features extracted using the block maxima approach are then used to fit the GEV distribution. As the number of samples increases i.e., $n \to \infty$, the distribution of these extreme values converges to

$$G_{\mu,\sigma,\xi}(x) = \exp\left(-\left[1+\xi\left(\frac{x-\mu}{\sigma}\right)\right]^{-\frac{1}{\xi}}\right) \tag{2}$$

where μ is the location parameter, σ is the scale parameter and ξ is the shape parameter. Typically, maximum likelihood (MLE) is used to estimate these parameters by fitting the GEV distribution on the actual data.

Detecting Outliers using EVT

Once the distribution of extremes is learned, our interest is to find the emergencies by using these distributions. In EVT, return period and return level [3] are used to find the values that exceed the preset threshold. Return period can be defined as the time needed for an extreme event to reappear, i.e., the average number of blocks between two exceedence of the corresponding value. Whereas return level is the extreme value associated with the return period, i.e., block maxima value that is expected to exceed only once in the corresponding return period. The return level η exceeded once in m time period (return period), is given as:

$$P(X > \eta) = 1/m \tag{3}$$

where $p = 1/m$ is the probability of the false alarm to occur once in every m periods of time. To calculate the return level and period, the estimates of quantiles of the block maxima distribution are found by inverting equation (1):

$$\eta = {G_{\mu,\sigma,\xi}}^{-1}(1-p) \tag{4}$$

For different ξ, this is given as:

$$\eta = \begin{cases} \mu - \frac{\sigma}{\xi}[1-\{-\log(1-p)\}^{-\xi}] & \text{if } \xi \neq 0 \\ \mu - \sigma\log\{-\log(1-p)\} & \text{if } \xi = 0 \end{cases}$$

where η is the return level or the value associated with the return period m demonstrating the false emergency alarm.

Evaluation

To evaluate the proposed approach, we use a combination of actual data collected from the University of Connecticut (UConn)'s campus Wi-Fi network, and simulation. The data from UConn's network is used to assess the applicability of our approach in the real world. The main limitation of using AP association events is similar to the reason why EVT is ideal for our purpose: the fact that there is only a limited number of emergencies during the period of our study. Therefore, to stress-test our approach we use simulation based on the real-world data that we collected. In the following subsections we present the results of our evaluation.

Detecting Real-World Emergencies

We use AP association data collected from UConn campus Wi-Fi network coupled with the ground truth provided by the fire department to evaluate the performance of our proposed approach. The dataset spans over 5 months, including two different semesters (Fall 2012 and Spring 2013). In collaboration with UConn's Fire Department, we also collected emergency event logs to which the Fire Department responded. Using the record of emergency events, we were able to identify 7 total emergencies (mostly false fire alarms) that caused building evacuations in 4 different buildings throughout the campus. Due to the limitation in the ground truth data, we mainly investigate whether our approach can detect emergency events at the building level. The main feature we use is the number of exits from all the APs in a building. As discussed in the previous section, due to the rare nature of emergencies, EVT is used to detect outliers. Features are extracted after every 15 minutes; using other time interval lengths is left as future work. Using a user-defined *return period* and *return level*, an appropriate threshold is learned, beyond which an event is classified as an emergency. One of the major advantages of EVT is that the same return period and level can be used across buildings to automatically extract the appropriate thresholds.

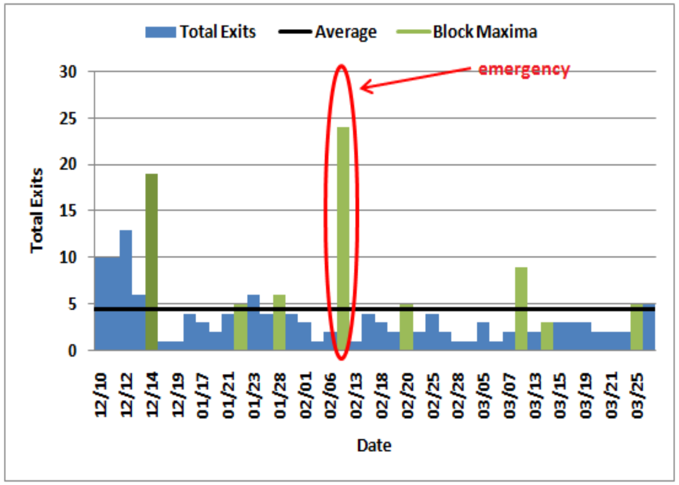

Figure 1: Example of an emergency and the block maxima for one of the buildings in our study.

Figure 1 shows an example of the number of exits in one of our buildings, where each of the bars corresponds to the number of exits in the same time window across days. The green bars indicate the maxima of each block that was used to extract the features for learning the GEV distribution.

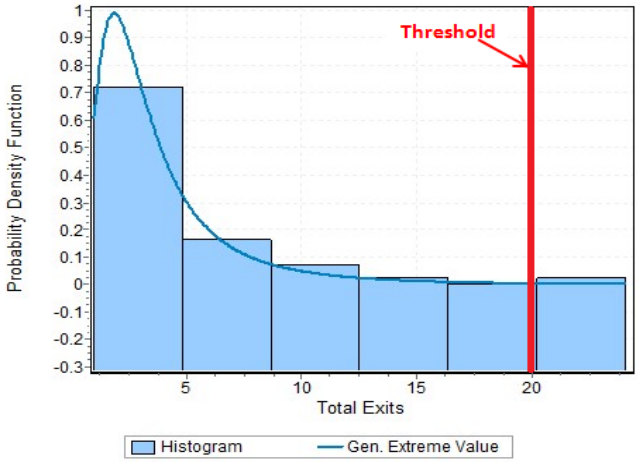

Figure 2: Example of extreme value distribution. It is learned for the data of Figure 1.

The values extracted using the block maxima approach are then used to fit a GVE distribution using MLE. Figure 2 shows the GVE that was learned for the data of Figure 1, as well as the threshold that is automatically inferred using equation (4). Using this threshold, emergencies can be detected by comparing it against the actual number of exits during this time window of the day.

To evaluate the performance of our approach using the real emergencies that we collected from the fire department's records we split the data into two sets: a training set containing all the instances of normal time windows, and a testing set containing both emergencies and normal behavior. Table 1 summarizes our results. The first row shows our emergency detection rate for each of the four buildings in our study, and the "False

Positives" indicates the number of times an emergency was erroneously detected.

Building	B_1	B_2	B_3	B_4
Detection Rate	$3/4$	$1/1$	$1/1$	$1/1$
False Positives	1	0	0	0

Table 1: Evaluation results using actual emergencies.

From Table 1 it appears that our approach performs particularly well for the few emergency samples that we were able to collect. More specifically, there is only a single missed emergency event out of four actual emergencies for building B_1, and a single false alarm again for the same building.

Using Simulation to Stress-Test our Approach
Since emergencies are rare, to further evaluate the performance of the proposed approach, we use trace-based simulation. Specifically, we simulate emergencies based on the original data collected from the Wi-Fi network. In order to simulate an emergency, a random date, time window and building is selected and the total number of devices present in the building is estimated using the AP association events. Subsequently, between 80% and 100% of the devices present in the building are randomly selected and are simulated to exit the building. Using this approach, we simulate 17 settings, each with 6 emergencies. Table 2 summarizes the results of our experiments. The "Detection Rate" column summarizes the number of simulated emergencies that were correctly detected for each building, and the "False Positives" indicates the number of false emergencies that were generated.

Setting	**Detection Rate**	**False Positives**
1	5/6	0
2	5/6	0
3	5/6	0
4	5/6	0
5	4/6	0
6	5/6	0
7	5/6	0
8	5/6	0
9	5/6	0
10	4/6	1
11	4/6	0
12	5/6	1
13	4/6	0
14	5/6	0
15	4/6	0
16	4/6	0
17	5/6	0
Summary	79/102	2

Table 2: Evaluation results using trace-driven simulation.

The overall detection rate of our approach is approximately 77.5%. This result is somewhat lower than what we anticipated, the main reason being that in some of our experiments, the "normal" number of people in the building is very low (e.g., later in the day, early in the morning etc.). This means that even in the case of an emergency the number of exits remains low. On the other hand, EVT appears to perform as expected with respect to the number of false alarms generated, causing only two false alarms across all buildings. Note that for this experiment, our goal was to demonstrate that EVT can be used to discover most of the simulated emergencies in each building, without generating a significant number of false alarms. We anticipate that by properly setting the return period and level we will be able to increase our detection rate but at a cost of a higher number of false alarms. Concerning our detection rate, even though in our preliminary experiments it appears to be rather significant, we expect that it can be greatly improved by considering more features and the temporal properties of the data. More specifically, in our current approach we only consider the number of exits from each building. Other features such as the number of people entering or present in the building could also prove useful for identifying outliers. Similarly, taking into account time information (e.g., day of the week) could help improve the accuracy of the learned GEV distributions.

Conclusion and Future Work

In this paper, we proposed a novel approach that locates emergencies in a campus using AP association data collected from the campus Wi-Fi network. The proposed method uses network information and traffic patterns to identify smartphones. Subsequently, the AP association events generated by smartphones are used to extract privacy-preserving features that describe the aggregate behavior of people at different areas of the campus for different time intervals of the day. Using the principles of EVT, an "extreme" distribution is extracted from the features, and is used to automatically learn a "threshold" above which an event can be classified as "emergency." To evaluate our approach we use data collected from our university's Wi-Fi network and emergency logs from its fire department over a period of 5 months. Our preliminary results demonstrate that our approach holds a lot of promise for identifying emergencies without generating a significant number of false alarms. We expect that a more thorough study using the existing and additional data from our university's network, will allow us to fine-tune our approach and improve our results. More specifically, we plan to explore the use of additional features in our models, as well as the use of multi-variate EVT [8].

As future work, we will pursue two main directions: (1) how to localize the initial area of an emergency, and (2) how to track the evolution of an emergency over space and time.

Acknowledgements

We would like to thank J. Farese, S. Maresca, and J. Pufahl of UConn University Information Technology Services and M. Makuch and J. Mancini of UConn Fire Department for their help with the project. The work was partially supported by National Science Foundation under Grant No. 1251962 and NSF CAREER award 0746841. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation. The authors would also like to thank the anonymous reviewers for their insightful comments.

References

[1] Afanasyev, M., Chen, T., Voelker, G. M., and Snoeren, A. C. Analysis of a mixed-use urban WiFi network: when metropolitan becomes Neapolitan. In *Proc. of IMC* (2008).

[2] Bamis, A., Lymberopoulos, D., Teixeira, T., and Savvides, A. The behaviorscope framework for enabling ambient assisted living. *Personal and Ubiquitous Computing 14*, 6 (2010), 473–487.

[3] C. Hor, S. Watson, D. I., and Majithia, S. Assessing load forecast uncertainty using extreme value theory.

[4] Chakrabarti, D. Autopart: Parameter-free graph partitioning and outlier detection. In *PKDD*, J.-F. Boulicaut, F. Esposito, F. Giannotti, and

D. Pedreschi, Eds., vol. 3202 of *Lecture Notes in Computer Science*, Springer (2004), 112–124.

[5] Chandra, R., Padhye, J., Wolman, A., and Zill, B. A location-based management system for enterprise wireless lans. NSDI (2007).

[6] Chen, X., Jin, R., Suh, K., Wang, B., and Wei, W. Network performance of smart mobile handhelds in a university campus wifi network. In *Proc. of ACM IMC* (November 2012).

[7] Cheng, J., Wong, S. H., Yang, H., and Lu, S. Smartsiren: virus detection and alert for smartphones. In *Proceedings of the 5th international conference on Mobile systems, applications and services*, MobiSys '07, ACM (New York, NY, USA, 2007), 258–271.

[8] Clifton, D. A., Hugueny, S., and Tarassenko, L. Novelty detection with multivariate extreme value statistics. *Journal of signal processing systems 65*, 3 (2011), 371–389.

[9] Debbabi, N., Kratz, M., Mboup, M., and El Asmi, S. Combining algebraic approach with extreme value theory for spike detection. In *Signal Processing Conference (EUSIPCO), 2012 Proceedings of the 20th European*, IEEE (2012), 1836–1840.

[10] Dick, A., and Brooks, M. Issues in automated visual surveillance. In *International Conference on Digital Image Computing: Techniques and Applications* (2003).

[11] Dick, A. R., and Brooks, M. J. Issues in automated visual surveillance. In *Proc. VIIth Digital Image* (2003), 195–204.

[12] Gember, A., Anand, A., and Akella, A. A comparative study of handheld and non-handheld traffic in campus WiFi networks. In *Proc. of PAM* (2011).

[13] Hsu, W.-j., Dutta, D., and Helmy, A. Mining behavioral groups in large wireless lans. In *Proceedings of the 13th annual ACM international conference on Mobile computing and networking*, MobiCom '07, ACM (New York, NY, USA, 2007), 338–341.

[14] Ko, T. A survey on behavior analysis in video surveillance for homeland security applications. In *Applied Imagery Pattern Recognition Workshop, 2008. AIPR'08. 37th IEEE*, IEEE (2008), 1–8.

[15] Kumar, U., and Helmy, A. Extract: mining social features from wlan traces–a gender-based case study. In *Proceedings of the 13th ACM international conference on Modeling, analysis, and simulation of wireless and mobile systems*, MSWIM '10, ACM (New York, NY, USA, 2010), 240–247.

[16] Kwan, M., and Lee, J. Emergency response after 9/11: the potential of real-time 3d gis for quick emergency response in micro-spatial environments. *Computers, Environment and Urban Systems 29*, 2 (2005), 93–113.

[17] Lakhina, A., Crovella, M., and Diot, C. Mining anomalies using traffic feature distributions. In *Proceedings of the 2005 conference on Applications, technologies, architectures, and protocols for computer communications*, SIGCOMM '05, ACM (New York, NY, USA, 2005), 217–228.

[18] Moghaddam, S., Helmy, A., Ranka, S., and Somaiya, M. Data-driven co-clustering model of internet usage in large mobile societies. In *MSWiM*, V. R. Syrotiuk, F. Alagz, B. Bensaou, and z. B. Akan, Eds., ACM (2010), 248–256.

[19] Nychis, G., Sekar, V., Andersen, D. G., Kim, H., and Zhang, H. An empirical evaluation of entropy-based traffic anomaly detection. In *Proceedings of the 8th ACM SIGCOMM conference on Internet measurement*, IMC '08, ACM (New York, NY, USA, 2008), 151–156.

[20] Papapanagiotou, I., Nahum, E. M., and Pappas, V. Configuring DHCP leases in the smartphone era. In *Proc. of IMC* (2012).

[21] Rauschert, I., Agrawal, P., Sharma, R., Fuhrmann, S., Brewer, I., and MacEachren, A. Designing a human-centered, multimodal gis interface to support emergency management. In *Proceedings of the 10th ACM international symposium on Advances in geographic information systems*, ACM (2002), 119–124.

[22] Smith, R. L. Extreme value theory. *Handbook of applicable mathematics 7* (1990), 437–471.

[23] Staniford-chen, S., Cheung, S., Crawford, R., Dilger, M., Frank, J., Hoagl, J., Levitt, K., Wee, C., Yip, R., and Zerkle, D. Grids - a graph based intrusion detection system for large networks. In *In Proceedings of the 19th National Information Systems Security Conference* (1996), 361–370.

[24] Turoff, M., Chumer, M., Van De Walle, B., and Yao, X. *The Design of a Dynamic Emergency Response Management Information System (DERMIS).* National Emergency Training Center, 2004.

[25] Zdziarski, E., Dunkel, N., and Rollo, J. *Campus Crisis Management: A Comprehensive Guide to Planning, Prevention, Response, and Recovery*. The Jossey-Bass higher and adult education series. Wiley, 2007.