Poster Abstract: WiFiMon: A Mobility Analytics Platform for Building Occupancy Monitoring and Contact Tracing Using WiFi Sensing

Emmanuel Cecchet*, Amrita Acharya†, Tergel Molom-Ochir*, Amee Trivedi*, Prashant Shenoy*
* University of Massachusetts, Amherst MA USA
†Smith College, Northampton MA USA
*{cecchet, tmolomochir, amee, shenoy}@cs.umass.edu †aacharya@smith.edu

ABSTRACT

With the current COVID-19 pandemic, contact tracing and building occupancy tracking are key components of re-opening policies and quickly containing virus outbreaks. WiFiMon is a network-centric contact tracing method that uses enterprise WiFi networks logs for tracking devices and inferring building occupancy and building contact tracing reports in office and campus settings.

We have built WiFiMon on an analytics platform that leverages the popular open source ELK stack (Elasticsearch, Logstash and Kibana) to process syslog data at large scale and offer users dashboards to quickly monitor building occupancy and generate contract tracing reports.

WiFiMon has been deployed at several UMass campuses gathering daily data from over 5,000 WiFi access points and tens of thousands of users. Our software is freely available for anyone to use and can be found at https://github.com/umassos/elastic-wifitrace.

CCS CONCEPTS

• Information systems ~ Mobile information processing systems

KEYWORDS

Wifi sensing, Contact tracing, Mobility analytics, ELK stack

1 Introduction

Mobility analysis can be an important approach for contact tracing and building occupancy monitoring for containment and mitigation of Covid 19. Analyzing human mobility can enable digital contact tracing which when combined with testing has emerged as an effective method of mitigation of Covid 19 [1]. Building occupancy monitoring can help with better planning and prevention by detecting concentration hotspots especially on university or enterprise campuses. In our previous work [2], we have developed a network-centric approach for phone-based contact tracing that relies on WiFi sensing [3]. In contrast to client-side approaches that depend on the use of Bluetooth and mobile applications, a network-centric approach does not require data collection to be performed on the device or apps to be downloaded by the user on the phone. Instead we collect system logs (syslogs) from WiFi access points (APs) that are readily available in enterprise WiFi networks and develop an analytics platform to infer user mobility patterns from chronological logs of their AP associations. Analytic queries over these mobility patterns are used to monitor occupancy levels in various parts of a building over time as well as to conduct digital contact tracing.

WiFiMon is an analytics platform built using the popular open source ELK stack (Elasticsearch, Logstash and Kibana) that is commonly used to process large amounts of syslog data in real time. As shown in Figure 1, the syslog files coming from the WiFi APs are processed by Logstash that takes care of the parsing and extracting the basic association/disassociation events that describe when a device MAC address joins or leaves an Access Points (AP). The location of APs in buildings can be used to figure out building occupancy (or density of devices around APs) and the proximity of users at any given time. All the processed data is stored in Elasticsearch, a distributed RESTful, JSON-based search engine.

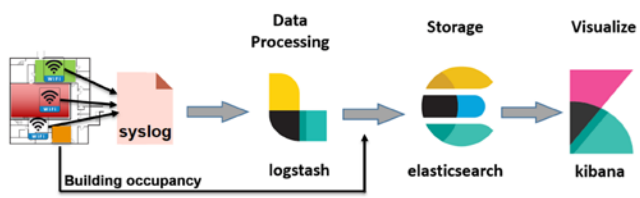

Figure 1. Overview of WiFiMon components

Kibana is the user facing interface that queries the data from ElasticSearch, aggregates and presents it in various forms to the user in dashboards. Dashboards can contain interactive tables, graphs, heatmaps and advanced visualizers to quickly zoom in a particular time slice or filter the data per location, user or AP.

The rest of this document is structured as follows: section 2 describes how we extract and store WiFi sensing information. Section 3 describes how building occupancy tracking is performed while section 4 presents our contact tracing technique. Section 5 concludes this paper.

SenSys '20, November 16–19, 2020, Virtual Event, Japan

ACM ISBN 978-1-4503-7590-0/20/11…$15.00
https://doi.org/10.1145/3384419.3430598

2 Extracting WiFi sensing data

Our current deployments rely on syslog data from HP Aruba APs. We wrote a specific Logstash parser that parses the Aruba message format and extracts the user device MAC address, AP information, association or disassociation event and the user information if this was a successful authentication event. We can parse either anonymized log files like on the UMass Amherst campus where all user and MAC address information is encrypted using a salted hash, or clear information when the software is running internally to the IT department like at the UMass Medical school campus.

Elasticsearch is not a typical SQL database, instead it stores data as JSON documents which is great to store unstructured data but makes some common tasks such as join operations with structured data a little bit harder to achieve. Therefore, some information has to be duplicated at load time to make it easier to aggregate later on. As such we have created ingest pipelines that merge information such as AP location in buildings and floors as well as maximum building occupancy data, with data from the log files. We also track session information (start, duration and end time between association and disassociation of each MAC address at each AP) on the fly in a Logstash aggregator, so that data becomes readily available in each JSON document and can be aggregated and queried directly by Kibana visualizers to be displayed in dashboards(see Figure 2).

Figure 2. Elastic WiFiMon Kibana main dashboard

3 Building occupancy tracking

Re-opening campuses during the Covid-19 pandemic oftentimes requires reduced building occupancy to be able to achieve social distancing guidelines. Each building occupancy information can be provided in a CSV file that is automatically processed by the ingest pipeline and computed on the fly as data is being loaded.

The WiFi sensing data we collect is aggregated by Elasticsearch queries to compute the number of devices at each access point at any given time. When the AP location is provided, it can be further aggregated to compare it to floor/building occupancy requirements. The results of these queries generate occupancy heatmaps per building, floor and AP per 5-minute period to quickly highlight hotspots throughout the day on campus.

An operator can quickly zoom in on the hot zones and analyze which parts of buildings are near or over maximum capacity throughout the day. Users can further plug in Kibana alerts which are recurrent queries that trigger actions if specific values are above certain thresholds.

4 Contact tracing

When a user has tested positive for Covid, we are looking for the other people the user has been in contact with. We provide a tool that generates a contact tracing report that contains the list of users who were connected to the same AP as a traced user for prolonged amount of time (current CDC guidelines are 15 minutes or longer contact periods). The tool calculates the contact time at each AP location and return a sorted list of users that have been the most exposed. We also extend the search to all other APs on the same floor of the same building as several APs might cover the same physical space depending on the layout. It is often possible in auditoriums or conference rooms to have 2 users in contact while being actually connected to different APs.

While the plain text contact reports can be used as is by a contact tracer, they can also be automatically loaded in Elasticsearch to be visualized in a contact tracing dashboard. Further alerts can be easily set to make sure that quarantined Covid patients are not coming in contact with the general campus population until their quarantine is completed.

Note that unlike Bluetooth approaches, WiFi sensing does not give an accurate indication of proximity. It also only works where WiFi coverage is available and it will not be able to follow users once they are out of range. However, WiFiMon works out of the box and does not require specific apps or specific interventions from users which has hampered the effectiveness of other approaches.

5 Conclusion

We have presented WiFiMon, a WiFi sensing technology that can be used for both building occupancy tracking and contact tracing. We leveraged the EKL stack to process WiFi AP syslogs in a scalable way. We have built several Kibana dashboards to help Covid task forces in their efforts from contact tracing to building occupancy analysis.

We have successfully deployed the software on UMass campuses with more than 5,000 APs on the Amherst campus and tens of thousands of users. We are expanding our support to other AP manufacturers as we help others deploy this technology. The software and documentation are freely available from our GitHub repository at https://github.com/umassos/elastic-wifitrace.

ACKNOWLEDGMENTS

This work has been funded by NASA grant NNX15AD92G and ARL grant W911NF-17-2-0196.

REFERENCES

[1] Joel Hellewell et al. 2020. Feasibility of controlling COVID-19 outbreaks by isolation of cases and contacts. The Lancet Global Health (2020)

[2] Amee Trivedi, Camellia Zakaria, Rajesh Balan, Prashant Shenoy, 2020. WiFiMon: Network-based Contact Tracing for Infectious Diseases Using Passive WiFi Sensing, arXiv:2005.12045.

[3] Dheryta Jaisinghani, et al.. 2018. Experiences & Challenges with Server-Side WiFi Indoor Localization Using Existing Infrastructure. In Proceedings of the 15th International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services, MobiQuitous New York City, NY, Henning Schulzrinne and Pan Li (Eds.). 226–235