Poster: Automatic Mass Power Outage Detection in Radio Access Networks

Milla Lintunen
Elisa Oy, Finland
University of Helsinki, Finland

Gopika Premsankar
University of Helsinki, Finland

Henri Tenhunen
Elisa Oy, Finland

Sasu Tarkoma
University of Helsinki, Finland

Ashwin Rao
University of Helsinki, Finland

ABSTRACT

Mobile devices are expected to be always connected, and this implies that the mobile network is able to quickly identify and address faults that impact service (for example, due to power outages). In this article, we present our approach for automatically detecting mass power outages. Our solution decreases the number of created trouble tickets in two mobile networks by 4.7% and 9.3%.

CCS CONCEPTS

• **Networks** → **Mobile networks**; • **Information systems** → **Data streaming**.

KEYWORDS

alarm correlation, mobile networks, spatio-temporal clustering, fault management

ACM Reference Format:
Milla Lintunen, Gopika Premsankar, Henri Tenhunen, Sasu Tarkoma, and Ashwin Rao. 2023. Poster: Automatic Mass Power Outage Detection in Radio Access Networks. In *The 21st Annual International Conference on Mobile Systems, Applications and Services (MobiSys '23), June 18–22, 2023, Helsinki, Finland.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3581791.3597365

1 INTRODUCTION

The increased proliferation of mobile devices in our daily lives requires them to be always connected, and this demands massive deployment of cellular base stations in the radio access network (RAN). The 5G mobile networks also rely on short-range and high-bandwidth mmWave technologies that require dense deployment of base stations.

Faults routinely occur in mobile networks for a wide range of reasons, for instance, due to disruptions in the electricity network (transmission or power failures), or configuration failures. These faults result in alarms generated by the network elements, which are processed in the Network Operations Center (NOC). The generated alarms contain several information fields, and the content in these fields is vendor-specific. Typically alarms contain a timestamp, alarm type, and other fields that include probable cause information. However, *alarms do not usually reveal the root cause of the failure*, and the NOC leverages the generated alarms to identify what happened in a specific network element. A typical life-cycle of an alarm and the various control plane elements traversed by an alarm generated in the RAN is presented in Figure 1.

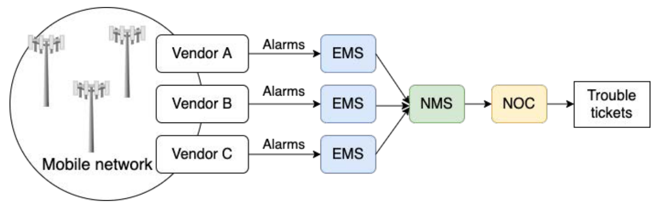

Figure 1: Lifecycle of an alarm. *Alarms traverse the Element Management System (EMS) and the Network Management System (NMS) before arriving at the Network Operations Center (NOC) where the network failures are handled.*

In the case of power failures, the NOC can receive alarms from multiple base stations. Identifying the root cause for alarms generated during mass power outages is non-trivial because a mobile network can rely on several electricity providers for powering the base stations; consequently, the electricity network topology is not correlated with the mobile network. Moreover, modeling and analyzing power failures is difficult, as they evolve over time in a complicated manner and spread randomly [4]. A failure to identify the root cause can result in multiple trouble tickets, and the investigation of separate tickets consumes time and increases costs for the network operator. For instance, field maintenance technicians can be sent simultaneously to multiple sites if problems are not detected as having the same root cause, and these technicians can end up traveling large distances which results in unnecessary costs.

In the following, we summarize our approach for detecting mass power outages that reduces the number of created trouble tickets. We also highlight key challenges in offering it as a service that can be included in systems presented in Figure 1. Our preliminary results show that we are able to detect mass outages in real time, and we can reduce the number of created trouble tickets and help reduce the costs of running the network.

2 SOLUTION OVERVIEW

Our goal is to identify mass outages by clustering the alarms according to temporal and spatial factors. We also believe that different types of mass outages, such as power and transmission outages, can be detected with the same method (or algorithm) by using different sets of fields from the alarms. Specifically, we argue that it

MobiSys '23, June 18–22, 2023, Helsinki, Finland

ACM ISBN 979-8-4007-0110-8/23/06.
https://doi.org/10.1145/3581791.3597365

is possible to cluster the alarms into groups and detect possible mass outages. We test our hypothesis using a dataset from a mobile operator with almost 2 million alarms gathered over a duration of a few months.

We leverage domain experts to filter power-related alarms indicating the base station site has lost power. We consider an alarm having a short duration as a flapping alarm, and in our case, if the site has lost its power for a longer duration (e.g., 10 minutes), we assume a permanent failure that requires corrective actions. In our dataset, we observe that only a small percentage (~8.5%) of power outage alarms are active for more than 10 minutes. Although alarms are raised within close time intervals of each other during a mass outage, a timestamp alone cannot reliably group the sites affected by the outage. We therefore consider different scenarios, with two parameters: the activity time of the alarm (1 to 10 minutes), and the time window in which the alarm arrives.

We analyze our dataset using Density-Based Spatial Clustering of Applications with Noise (DBSCAN) with different parameters (epsilon or distance measure, and number of points). Our analysis indicates that the choice of DBSCAN parameters depends on the density of base station sites. Figure 2 shows that a mass outage is detected correctly in a rural area (with fewer base station sites), whereas a mass outage is incorrectly identified in an urban area with the same parameters. We therefore propose an algorithm that groups alarms into clusters based on the estimated density of sites for each incoming alarm in an online manner, similar to SOStream [1]. We use the k-Nearest Neighbours (kNN) to estimate the density and then set the distance up to which neighbors of a site must be considered; specifically, we first calculate the average distance of a site to k (k is empirically determined and validated with domain experts) nearest sites. We then introduce a *scaling factor* to increase the size of the considered radius for detecting neighboring sites for each incoming alarm's site. We use these density estimates to cluster and group incoming alarms in an online manner; details of the algorithm are available in our report [2].

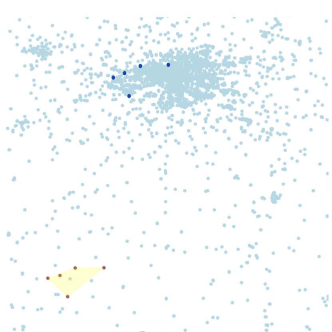

Figure 2: Issues with using a fixed value for epsilon (distance). *A mass outage, highlighted in yellow, is correctly identified in a rural area. The same algorithm incorrectly identifies a mass outage, marked with dark blue circles in the figure, in an urban area with a higher density of sites.*

We implemented our algorithm in a Mass Outage Detection Service to be used in a Virtual NOC at Elisa Polystar [3]. Our service uses Apache Kafka for streaming the alarms, PostgreSQL database as a data store, and PostGIS extension to query geographical information. We evaluated our detection service with a synthetic dataset comprising 10000 sites spread within Denmark. We generated 64 power alarms that contained the information of the site name and alarm occurrence time, and we chose the site and occurrence time for 35 of these alarms so that they form 5 different mass outage clusters.

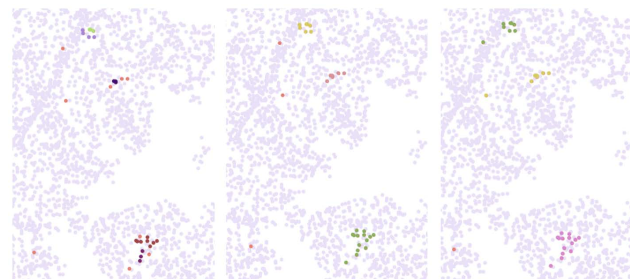

Figure 3: Impact of the scaling factor. *Different scaling factors (left: 1, middle: 3, right: 5) with kNN=10. Gray points represent sites without active power alarms. Salmon-colored points demonstrate outliers. Mass outage clusters are identified by similarly colored points. A small (1) and high scaling factor (5) include outliers to the clusters.*

We explored the impact of K and different scaling factors; Figure 3 illustrates a sample result. An extensive analysis of different parameter settings is presented in our report [2]. We also evaluated our algorithms in two real networks and found that detecting the mass outages in real-time resulted in a 4.7% and 9.3% reduction of created trouble tickets.

3 CONCLUDING REMARKS

We proposed a mass outage detection service that clusters power-related alarms in the RAN in real-time. Our preliminary results are promising, and we aim to integrate additional data sources in detecting mass outages, and further analyze the threshold settings and parameters of the algorithm. We also plan to evaluate the efficacy of our detection service in more networks with different densities of sites.

4 ACKNOWLEDGMENTS

This was work was supported in part by Academy of Finland under grants 345008 and 338854, and by Business Finland Neural Pub/Sub for 6G project.

REFERENCES

[1] Charlie Isaksson, Margaret H Dunham, and Michael Hahsler. 2012. SOStream: Self organizing density-based clustering over data stream. In *Machine Learning and Data Mining in Pattern Recognition (MLDM)*.

[2] Milla Lintunen. 2023. *Automatic Detection of Mass Outages in Radio Access Networks.* Master's thesis. University of Helsinki.

[3] K. Valtari and E. Vesterinen. 2020. Automate the Network Operation Center. Whitepaper.

[4] Lei Zhong, Kiyoshi Takano, Fangzhou Jiang, Xiaoyan Wang, Yusheng Ji, and Shigeki Yamada. 2016. Spatio-temporal data-driven analysis of mobile network availability during natural disasters. In *International Conference on Information and Communication Technologies for Disaster Management (ICT-DM)*.