Predicting Outages in Radio Networks with Alarm Data

Dan Kushnir
Bell Laboratories, Nokia
Murray Hill, NJ 07974, USA
dan.kushnir@nokia-bell-labs.com

Gautam Gohil
Mobile Networks Quality, Nokia
Bengaluru, India
Gautam.gohil@nokia.com

Zulfiquar Sayeed
Bell Laboratories, Nokia
Murray Hill, NJ 07974, USA
Zulfiquar.Sayeed@nokia.com

Huseyin Uzunalioglu
Bell Laboratories, Nokia
Murray Hill, NJ 07974, USA
huseyin.uzunalioglu@nokia.com

I. Introduction

Modern cellular networks are complex systems offering a wide range of services and present challenges in detecting anomalous events when they do occur. The networks are engineered for high reliability and, hence, the data from these networks is predominantly normal with a small proportion being anomalous. From an operations perspective, it is important to detect these anomalies in a timely manner in order to mitigate them and preclude the occurrence of major failure events. In telecommunications diverse set of data such as KPIs, logs and alarms are generated to monitor the health and stability of the network element and the services carried over it [1] [2] [3] [4].

While network operators monitor KPI and alarms, they are hampered by two issues:

- In many cases, due to their large number, all alarms cannot be monitored in real time, and as a result some alarms are flagged as critical and analyzed while others are dismissed as nuisance alarms even though a change of arrival rate in those ”nuisance” alarms could indicate an upcoming outage.
- There is no end-to-end system view or model providing a clear relationship between an alarm and a BSC failure - and corrective action is not taken prior to the incident.

The fault tolerant mechanisms in the network correct most errors, thereby masking any short-duration anomalies that could have occurred. At times, these anomalies may be a precursor to a larger failure in the system. It is therefore important to detect such ”needle in the haystack” anomalies in a timely fashion to identify vulnerabilities and take corrective measures as necessary.

Our goal in this paper is to anticipate upcoming outages that have vast impact on the network service *days* before they occur. We propose a statistically rigorous methodology for anomaly detection in alarm generation frequencies. Our approach is based on the assumptions that an anomalous increase in frequency serves as a precursor to a failure. This novel methodology can be applied to any network element that generates alarms, as well as to other systems beyond networks. The anomaly detection tool (ADT) builds a statistical model for each elements alarm-type generation process by using a probabilistic Poisson process. The Poisson model allows to profile alarm generation rates during normal and healthy periods. Deviations in real-time from the learned rates are flagged as anomalies, from which those with an increased alarm frequency are of interest. ADT ranks the propensity of each element to go into outage by an anomaly scoring system which allows to score network elements with highest propensity to fail.

We focus our tool’s application on a test case of alarms generated by Base Station Controllers in a 2G network. Our evaluation demonstrates that outages can be detected by ADT at least one day before outage occurrence, thus, allowing operations teams to take proactive steps to cure problems and prevent outages.

II. Methodology

In this section we expand on our algorithm for learning anomalous behavior that potentially leads to outage. For simplicity, we address the network element of interest as a Base Station Controller (BSC), however, other alarm generating elements may be considered.

A. Preliminaries

Data: We consider our data set as a dynamic time series over time intervals $T_n = t_1, ..., t_n$, where t_i is a time interval of fixed length $\triangle$ for all i. We define an alarm-counter vector $A(B)^q = (a^q_1, ..., a^q_n)$ such that an alarm of type q has had a^q_i occurrences during the time period t_i for BSC B. If m dictates the number of alarm-types that are under observation we have that $A(B_j) \in \mathbb{R}^{m \times n}$ is an $m \times n$ matrix for BSC B_j. To facilitate notation, a set of particular purpose, such as a training set of alarm occurrence observations for BSC B over a time period T_n, is denoted by $A_{train}(B)$.

978-1-5386-2542-2/18/$31.00 © 2018 IEEE

Poisson Process Modelling: The Poisson point process is related to the Poisson distribution [5], which implies that the probability of a Poisson random variable K being equal to k is given by:

$$P\{K = k\} = \frac{\lambda^k}{k!} e^{-\lambda}, \tag{1}$$

where λ is the Poisson distribution rate parameter. (1) is often used to model the number of times an event occurs in an interval of time. We observe that alarms occur in a network element during healthy times with a certain characteristic frequency. However, before outage this frequency increases due to malfunctions. We therefore use the Poisson process to model the healthy frequency and attempt to detect statistically significant deviation from the health-characteristic frequency as a precursor for an outage.

B. Training

Formally, we are interested in learning the rate $\lambda^q(B)$ for alarm type q at BSC B during healthy periods, i.e. when no outages occur. One should therefore expect that over a long enough period, in which no outage has occurred on that BSC, the average rate $\bar{\lambda}^q(B)$ over the set T_n should reflect a rate that characterizes normal alarm generation. To this end, we can compute the mean rate over the period $T = t_1, ..., t_n$ as

$$\bar{\lambda}^q(B) = \frac{1}{n} \sum_i a_i^q(B). \tag{2}$$

$\bar{\lambda}^q(B)$ is used to estimate the probability of an alarm of type q to be triggered k times at BSC B during time interval of length $\triangle$, according to a Poisson process with $\bar{\lambda}^q(B)$:

$$P\{k\ alarm\ events\ occured\ in\ \triangle\} = \frac{(\bar{\lambda}^q(B))^k}{k!} e^{-\bar{\lambda}^q(B)}. \tag{3}$$

C. Anomaly Detection

We utilize the Poisson model to compute a probability score for each observed alarm-rate being abnormally high. Specifically, we assess the probability that an observed rate λ_{test} at time t_{test} is at the right tail of the distribution of alarm frequency with mean healthy baseline rate $\bar{\lambda}$, by examining the Cumulative Distribution Function (CDF) $P(X \leq \lambda_{test})$. If $P_{\bar{\lambda}^q}(X \leq \lambda_{test})$ is high, then the observed rate λ_{test} is considered anomalous and much larger than the probability derived by assuming a rate of $\bar{\lambda}^q$. A predefined threshold α can be set to flag anomalous rates. We typically use $\alpha = 0.95$. An indicator function for an anomaly of an alarm type q with referral to baseline mean alarm rate $\bar{\lambda}^q$ can be defined as

$$I_\alpha(\lambda^q(B)) = \begin{cases} 1 & if\ P_{\bar{\lambda}^q}(X \leq \lambda_{test}^q) > \alpha \\ 0 & otherwise \end{cases}$$

D. Scoring and ranking likelihood to fail

To this end, we have the probabilistic apparatus to generate a scoring and ranking mechanism for the likelihood of a BSC to fail. Clearly, a BSC (or any network element) that has a high score should be subject to at least diagnostics, if not proactive action taken. One basic mode of operation would be counting the number of anomalies that occurred per a BSC. We therefore define the following score as $\Gamma_w(B) = \sum_q w_q I_\alpha(\lambda_{test}^q(B))$, where w_q is a weight corresponding to alarm type q. Alternatively, one can use the actual probability $\Gamma_P(B) = \sum_q P_{\bar{\lambda}^q}(X \leq \lambda_{test}^q)$, instead of the indicator function.

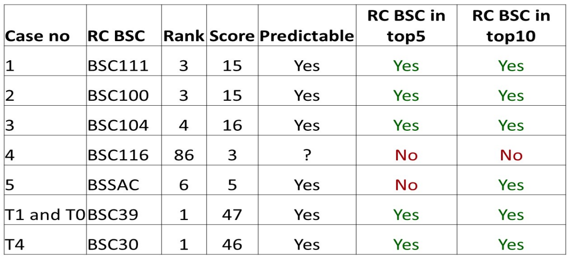

Case no	RC BSC	Rank	Score	Predictable	RC BSC in top5	RC BSC in top10
1	BSC111	3	15	Yes	Yes	Yes
2	BSC100	3	15	Yes	Yes	Yes
3	BSC104	4	16	Yes	Yes	Yes
4	BSC116	86	3	?	No	No
5	BSSAC	6	5	Yes	No	Yes
T1 and T0	BSC39	1	47	Yes	Yes	Yes
T4	BSC30	1	46	Yes	Yes	Yes

Fig. 1. Summary of ADT outage prediction results.

Randomization. Healthy baseline deviations that are due to temporary\periodic\rare occurring network events may cause healthy baseline rate computations that are not representative. Therefore, to increase the robustness of our scoring engine, we examine a set of mean scores over random selection of baseline time periods during training\baseline-learning stage: $\Gamma_1(B), ..., \Gamma_r(B)$ generated according to different mean rates $\bar{\lambda}_1^q(B), ..., \bar{\lambda}_r^q(B)$ for all q computed according to eq. (2) for uniformly at random selected time periods $T_1, ..., T_r$ ($T_i = t_{1,i}, ..., t_{n,i}$).

Next, we generate a mean score for BSC B as $\bar{\Gamma}(B) = \frac{1}{r}\sum_{i=1}^{r} \Gamma_i$, and generate an average ranked list $\Psi = i_1, ..., i_m$ of scores $\bar{\Gamma}(B_1), ..., \bar{\Gamma}(B_m)$, such that $\bar{\Gamma}(B_{i_1}) \geq \bar{\Gamma}(B_{i_2}) \geq ... \geq \bar{\Gamma}(B_{i_m})$ given m examined BSCs in the network.

III. Evaluation

In table 1 we summarize ADT's performance for a selected set of outage events occurring in a service provider network that includes 150 BSCs. The table includes information on the root causing BSC (RC BSC) scores, ranks, as well as the ADT's inclusion of the RC BSC in the top-5 and top-10 ranks *at least 1 day before the outage occurred* during the 5 days preceding to the outage. As seen, in most cases the tool predicts the BSC among the top-5 rank list out of 150 BSCs.

References

[1] G.F. Ciocarlie, U.Lindqvist, S.Nováczki, and H.Sanneck. Detecting anomalies in cellular networks using an ensemble method. In *Network and Service Management (CNSM)*, pages 171–174. IEEE, 2013.

[2] V. K. Gurbani, D. Kushnir, V. Mendiratta, C. Phadke, E. Falk, and R. State. Detecting and predicting outages in mobile networks with log data. In *2017 IEEE International Conference on Communications (ICC)*, pages 1–7, 2017.

[3] C. S. Hood and C. Ji. Proactive network-fault detection [telecommunications]. *IEEE Transactions on Reliability*, 46(3):333–341, 1997.

[4] C. Lim, N. Singh, and S. Yajnik. A log mining approach to failure analysis of enterprise telephony systems. In *2008 IEEE International Conference on Dependable Systems and Networks With FTCS and DCC (DSN)*, pages 398–403, June 2008.

[5] J. Kingman. *Poisson Processes*. Clarendon Press, 1992.