Correlation-Based Cell Degradation Detection for Operational Fault Detection in Cellular Wireless Base-Stations

Muhammad Zeeshan Asghar[1], Richard Fehlmann[2], and Tapani Ristaniemi[3]

[1] Magister Solutions Ltd, Hannikaisenkatu 41, Jyväskylä, Finland
[2] Nokia Siemens Network Research, Linnoitustie 6, Espoo, Finland
[3] Department of Mathematical Information Technology,
University of Jyvaskyla, Jyvaskyla, Finland
muhammad.asghar@magister.fi, richard.fehlmann@nsn.com,
tapani.ristaniemi@jyu.fi

Abstract. The management and troubleshooting of faults in mobile radio networks are challenging as the complexity of radio networks is increasing. A proactive approach to system failures is needed to reduce the number of outages and to reduce the duration of outages in the operational network in order to meet operator's requirements on network availability, robustness, coverage, capacity and service quality. Automation is needed to protect the operational expenses of t he network. Through a good performance of the network element and a low failure probability the network can operate more efficiently reducing the necessity for equipment investments. We present a new method that utilizes the correlation between two cells as a means to detect degradations in cells. Reducing false alarms is also an important objective of fault management systems as false alarms result in distractions that eventually lead to additional cost. Our algorithm is on the one hand capable to identify degraded cells and on the other hand able to reduce the possibility of false alarms.

Keywords: Mobile Networks, Fault Managements, Degradation Detection, Correlation, Operational Expenditures (OPEX), Capital Expenditures (CAPEX), Long Term Evolution (LTE), Self-Organizing Networks(SON).

1 Introduction

The huge amount of network elements in the mobile communication networks and consequently handling of huge amount of measurements recorded at each base station is a great challenge. The tasks of operation and maintenance of mobile cellular networks are not only vulnerable to errors as wireless network is complex to handle but also requires a huge amount of human resource to monitor and execute these tasks. A lot of effort is being put on this area of operations and maintenance of mobile communication systems that aims at simplifying the management of cellular network on the one hand and improving the efficiency by introducing automation, on the other hand. The fundamental requirement to tackle this challenge is to automate the

D. Pesch et al. (Eds.): MONAMI 2013, LNICST 125, pp. 83–93, 2013.

operations and maintenance functionalities. This triggered the concept of self-healing networks. Self-healing is an important domain of Self-Organizing Networks (SON). Other major domains of SON enabled networks are self-configuration and self-optimization.

Self-healing enabled cellular networks are generally defined as the wireless cellular networks where the tasks of troubleshooting including detection, diagnosis, corrective actions are largely automated but the operator will have final control over the decisions. The self-healing use cases defined by 3GPP [1] are given as: Self-Recovery of NE Software, Self-Healing of Board Faults, Cell Outage Detection, Cell Outage Recovery, Cell Outage Compensation and Return from Cell Outage Compensation. Although 3GPP use cases are focused on "cell outage", we adopt a more general concept of "cell degradation" which refers to the case where the actual performance of the cell in handling traffic is significantly lower as it is supposed to be. Degradations may not be measured directly as they do not necessarily trigger alarms. Degradations can be classified in terms of their severity i.e. from worse performance to complete outage. Usually a faulty cell starts degradation in prior to go in outage state. It is not enough to just detect the faults in a timely manner, but it is equally important to detect the degradations of the performance of cells (sectors) before real failures occur so that counter measures can be taken in time.

Recently two books are published on "Self-organizing networks" [2][3]. The class of 'Operational Fault Detection' OFD algorithms are introduced by Cheung [4]. The OFD algorithm analyze performance indicators detect fault signatures without the need for operators to manually set thresholds. The profile of the system is built by either looking at its earlier behavior or comparing it to similar systems. The correlation-based algorithm uses the correlations of cells (sectors) within a geographical neighborhood. It is assumed that there exists an appreciable level of correlation between neighboring cells. The same OFD approach is followed by a statistical hypothesis test framework for determining faults [5]. A method to detect coverage and dominance problems and identify interferers in WCDMA networks is introduced in [6]. Signaling messages exchanged through the radio interface are used to calculate certain metrics for every cell during normal network operations reflecting real traffic distributions and geographical user locations. Competitive neural algorithms are used for fault detection and diagnosis in 3G cellular networks in [7]. Another cell outage detection algorithm based on the neighbor cell list reporting of mobile terminals is introduced by Mueller in [8] . A framework is presented in [9] using Minimization of Drive Tests (MDT) databases, for detecting sleeping base stations, network outage and dominance areas in a cognitive and self-organizing way. Diffusion maps for reduction of high dimensionality and nearest neighbor classification methods were used. An experimental system for comprehensive testing of SON use cases is presented in [10]. A self-healing framework for 3GPP LTE networks is presented in [11] where detection and compensation of cell outages are evaluated in a realistic environment.

This paper is organized as follows: Section 2 describes the proposed algorithm for degradation detection. Section 3 discusses the procedure to choose the well-correlated cell pairs /comparing cells. Section 4 shows the investigations with the data recorded at a real 3G network. In this paper we do not give precise values of the performance metrics rather we address the vague terminologies in a better and clear way using empirical results.

2 The Proposed Algorithm

In our approach we exploited the idea that there are many cells in the network coverage area having similar behavior irrespective of their geographical locations. The terms "faultless situation", "load" and "appreciable level of correlation" will be used in the following sections, therefore, we give definitions of these terms here as follows.

The term "faultless situations" means the cells are functioning normally, i.e. they handle the traffic as they are supposed to and have no interior problems. Such cells are also called healthy cells. Although there might always be minor problems in the cell, for the sake of simplicity we may ignore them as long as they do not have significant impact on the performance of the cell.

"Load" can be represented by several key performance indicators. Some of the load related key performance indicators (KPIs) include traffic demand by users and others not. For example the KPI "throughput" is highly dependent on user traffic demands while "number of active users" is not. As the first one includes a higher degree of randomness, it is less usable in monitoring the correlation level of the load of the cell and so the number of active users in a cell is better suited to be used as a load level KPI to monitor the correlation levels.

"Appreciable level of correlation between two cells" refers to the fact that the relationship between the two cells is strong enough in faultless situations such that a diminishing of this correlation could be well noticed. We consider a correlation level as appreciable if for two cells whose correlation stay above that level in normal conditions a degradation of one of the cells is clearly visible as the correlation values start to drop then. It requires an effort to determine that level in a real network. However, network operators can identify a good value for appreciable correlation levels based on their field experience. Based on our empirical results we identified the value 0.5 as a good value to be used as appreciable level of correlation between two cells but this might be different in other networks. Two cells having an appreciable level of correlation, we call "well correlated cells". At this point we note that such cells need not to be adjacent cells.

Let us consider a KPI or a function of KPIs expressing the load handled at a cell, e.g. the number of active users allocated to the cell or the transferred downlink load. In regular time steps these values are recorded, let us denote the values for the target cell by x_1, x_2, …, x_n and for a comparing cell by y_1, y_2, …, y_n. The correlation coefficient r (X,Y) is given below where the vectors X and Y consist of the components $x_1, x_2, \ldots, x_n$, $y_1, y_2, \ldots, y_n$, respectively. Here, the sum is taken over the sliding window size n:

$$r(X,Y) = \frac{\sum_{i=1}^{n}(x_i - \overline{x})(y_i - \overline{y})}{\sqrt{\sum_{i=1}^{n}(x_i - \overline{x})^2}\sqrt{\sum_{i=1}^{n}(y_i - \overline{y})^2}} \quad (1)$$

The above mentioned formula is used to calculate the correlation values of each pair of cells. In a first step we have to choose cells where these correlation values consistently stay above some reasonable value, e.g. 0.5 as we proposed above. These chosen cells are fixed and referred to as the comparing cells. The appreciable level of correlation can be defined differently by the operator as stated above. The correlation threshold does not need to be the same for all networks.

We assume now that for each target cell a set of comparing cells has been determined. The correlation coefficients with each chosen comparing cell are now monitored. If a target cell's correlation coefficient with all its chosen comparing cells drops below the predefined reasonable bound there is a good reason to assume that a performance degradation of the target cell has started and a suspicion is raised therefore. At least if there is more than one comparing cell showing this behavior this is a good reason for an alarm. If there is only one such cell some additional cell performance monitoring is needed in order to determine if a cell is degrading and if yes which one. Such additional performance monitoring includes tracking KPIs of the cells and checking if they pass certain thresholds that operators have set in the traditional way. In any case it can be concluded that two or three comparing cells suffice to make a decision if an alarm shall be raised because if the target cell is degrading then it will be visible in all the correlation levels with comparing cells except for very unlikely exceptional situations. So if two or three comparing cells show decrease in correlation below that appreciable level it can be concluded that the target cell is degrading. On the other hand if the target cell is healthy the degradation levels with two or three comparing cell stay above that appreciable level and it can be concluded that the target cell is healthy without checking further comparing cells.

The question of precisely how many comparing cells should be required to detect the degradation should be answered by the operator. The question is how many comparing cells are needed in order for the operator to decide if a target cell needs some recovery measures.

We used a flexible critical time duration that has to be computed every time we utilize a fixed critical duration in connection with the sliding window. In a state of suspicion the target cell is locked. The monitoring follows more steps to see if this drop below the appreciable level is temporary only. The algorithm now waits for more steps. If during a few following steps the correlation of the target cell with its comparing cell(s) does not rise above the predefined level then an alarm will be triggered. Based on empirical results we have seen that the correlation between the comparing cell pair becomes stable as we increase the window length. Using this idea, we define the condition that must be met before triggering an alarm.

3 Investigations with Real Network Data

We had the opportunity to get access to a real 3G network and observe its performance data over a period of time. The observed network consists of thousands of cells and is located in a European city. The granularity of the performance data is one hour, thus each sample corresponds to a one hour measurement. In what follows we present a couple of typical correlation values statistics. They have been collected over a period of two weeks. In regular steps KPI values were measured and correlation coefficients calculated. The window size was 64 steps. The KPI represented the number of active users in the cell.

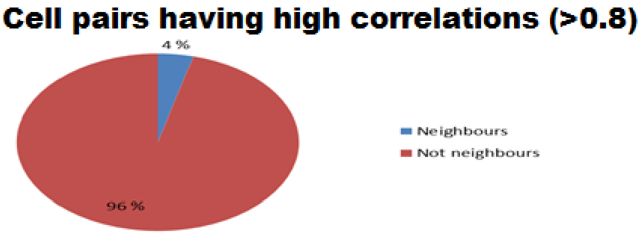

Fig. 1. Cell pairs having high correlations (>0.8)

The empirical results in Figure 1 show that the assumption of having high correlation with the geographical neighbors is not always correct. In the available data, out of all the cell pairs showing high correlation only 4% are neighbors whereas about 96% are cell pairs that are not geographical neighbors. Thus it proves that there are many more cells with high correlation than just the ones in the geographical neighborhood. For detection we can take all the available cells into account where the similarities in the performance with the target cell are reasonably high. It is up to the operator to decide how long to observe this similarity among the cells.

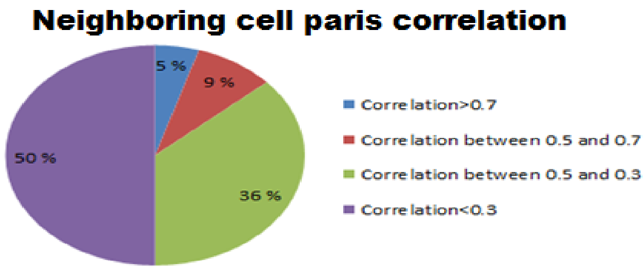

Fig. 2. Correlation of neighboring cell pairs

Figure 2 presents the statistics of the correlation values of all neighboring cell pairs. As can be seen, in spite of showing high resemble in the behavior over time the cells being geographical nearby show often low correlation. In this performance data collected from this real operational network, only 5 percentage of the cell pairs that are geographically nearby show high correlation and 14% show for our purposes still a good correlation level (above 0.5). So, one out of seven cells has an appreciable level of correlation. On the other hand 86 percentage of the cell pairs that are geographically neighbors, exhibit low correlation. This data shows that it is not reliable only to base the decision on the cell pairs that are neighbors to each other. However, two or three comparing cells for a target cell are enough for the algorithm as will be explained in section 5. Therefore, it can be expected that enough good correlated cells can be found among the closest 20 cells. There could be networks where relying on the geographical neighborhood gives good results. But this is not true always.

3.1 Degradation Types

As degradations are rare events and did not occur in our observed real data, we have introduced artificial errors in the real data. In order to produce degradation we have reduced the number of active users in the data as they would have been in case of a healthy cell. There are two types of degradation that we have introduced for the evaluation of the proposed methodology.

- Slow degradation
- Fast degradation

We introduced the slow degradation in the data collected from operational network by decreasing the number of active users in samples over time. Each step corresponds to a one hour measurement and thus the number of active users was decreased in every hour. Fast degradation indicates severe malfunction caused by the failing hardware component or software entity of the base station. We introduced fast degradation in the data collected from operational network by decreasing the number of active users to zero instantly.

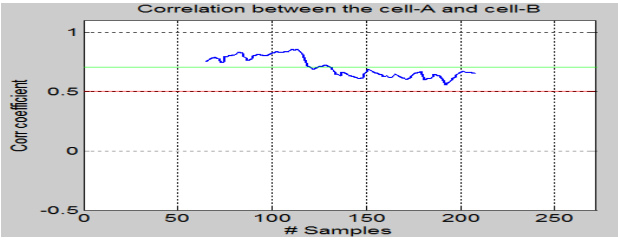

Fig. 3. Correlation between two cells based on the number of active users

Figure 3 presents a typical correlation value track represented by blue color. The green line shows the mean of the correlation curve and red line gives the threshold or the critical level of correlation. As can be seen the correlation curve starts at 65 (because the window size is 64) and show some variation of the correlation coefficient over time. In this example, however, these values stay consistently above

0.5 and are therefore useful for monitoring a possible degradation. The correlation value between cell-A and cell-B remains higher than the critical level in healthy conditions. The curve then ends at step 205 where monitoring had ended.

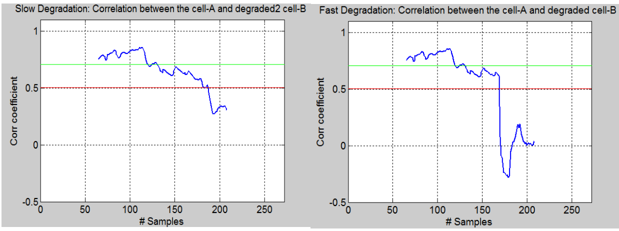

Fig. 4. a)Degradation type I: Cell-B degraded b) Degradation type II: Cell-B degraded

Figure 4 depicts the same coefficient as Figure 3 but with an artificially introduced slow degradation of the target cell (curve on the left) and a fast degradation (curve on the right) that starts at step 168. The slow degradation continues stepwise and as can be seen the curve falls down slowly. Such a type of degradation is easily spotted, however with a latency that depends on the window size.

In the fast degradation the number of active users in the cell is reduced to 0 all at once. As can be seen such a type of degradation is spotted fast and easily.

3.2 Window Size Analysis

There is a trade-off between the reaction time and accuracy of detection. For high reaction time less number of samples will be utilized whereas for accuracy it is better to have more samples to get the statistical confidence. An optimal number is required that satisfies both requirements. For instance, in our analysis, we proposed at least a two days period for an accurate detection on one hand and a quicker response to the network changes on the other hand. However, it is not easy to find a global optimum number that is valid for all cells in the network. For this reason, we emphasized the need of different observation time for each cell pair.

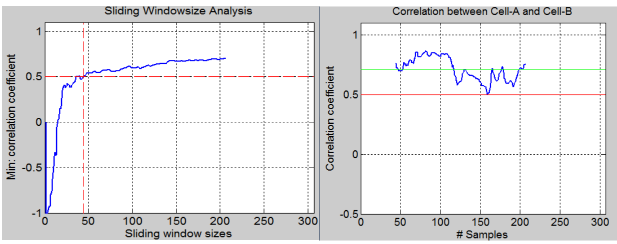

Fig. 5. a)Minimum Correlation Vs Window Size b) Correlation based on the window size 44

In the curve on the left in Figure 5, the optimal window size is depicted for an arbitrarily selected well correlated cell pair. In this example it turns out that 44 is the optimal window size meaning large enough for giving an accurate correlation value and as small as possible guaranteeing the smallest possible reaction time. As we further increase the window size the correlation value between the cell pairs stays above a certain level. This is because of the fact that by increasing the window size more samples included in the window that strengthen the accuracy of the correlation value between the cells. Although the further increase in the window size yields high correlation and solid results the reaction time will be decreased and this results in a slower detection process.

The curve on the right presents a correlation coefficient value track represented by blue color, with optimal sliding window size i.e. 44, as depicted in figure 5. The green line shows the mean of the correlation curve and red line gives the threshold or the critical level of correlation. As can be seen the correlation curve starts at 45 (because the window size is 44) and show some variation of the correlation coefficient over time, however, these values stay consistently above 0.5 and are therefore useful for monitoring a possible degradation.

4 Choice for Comparing Cells

In real networks user behavior varies strongly depending on the time of day, the geographic location and other factors. Although it might be that cells located near to each other would experience similar environmental conditions this is not always true as shown in the preceding section. It is also possible that some cells perform similar to other cells independent of their geographic location. This might be caused due to a similar kind of user behavior depending mainly on the time of the day. It is vital to choose the right cells as comparing cells. This is done in an initial cell pair selection process. In this section we describe this initial process.

In deciding how to choose comparing cells to a given target cell we follow two principles or requirements:

1. A degradation of the performance of a cell shall be detected
2. False alarms shall be avoided as good as possible

The first requirement implies that the correlation level between the target cell and its chosen comparing cell needs to remain high during faultless conditions, i.e. above the appreciable level of correlation. For a cell pair whose load correlation level stays above such a level in normal conditions it can be expected that if one of the cells starts to deteriorate this correlation levels will start to drop soon below that reasonable level. We have seen this in the preceding section where a degradation of the cell could easily be spotted as the monitored correlation level started to drop as soon as the degradation sets in. This suggests to observe all the cell pairs and to pick the ones whose correlation levels stay high.

However, this is very costly as the number of pairs grows like the square of the number of cells in question. It is therefore advisable to start with each target cell and a smaller group of cells to be monitored as potential comparing cells.

We have seen that the cell degradation implies a drop in the correlation level. But the converse is not true. If the target cell experiences a drop of the load correlation level with a comparing cell chosen as explained above then this is only one indication that the target cell is degrading. But there could be other reasons as well that the correlation level is dropping.

This information has to be combined with further information before the cell performance degradation can be declared. This can be done by looking at the correlation level with other chosen comparing cells or by taking traditional cell performance monitoring indicators into account. Only if several indicators have shown that the target cell is giving signs of deterioration it can be logically concluded that the target cell is degrading and an alarm should be raised.

We illustrate this by an example. Given a target cell let us assume we have two cells (green in figure 6) that can be paired with the target cell (purple) as both have correlation levels to the target cell above the critical level. Let us further assume that one of the cells is near the target cell and the other far away. The question is if there is an advantage by choosing one or the other.

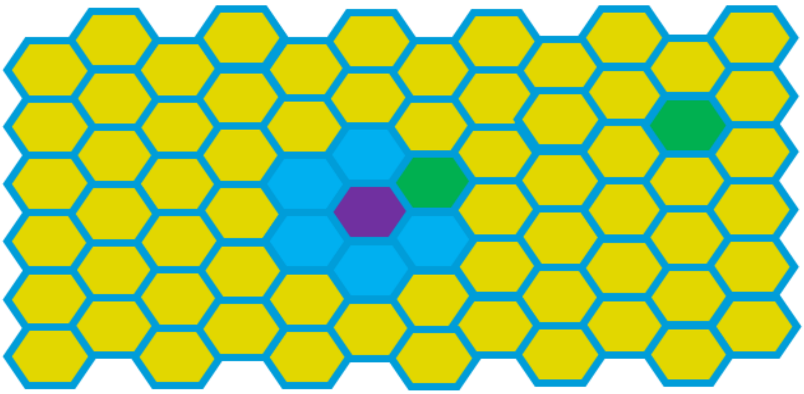

Fig. 6. Example, near and far away cells with good correlation in green

Let us now assume such a scenario as explained above where a peak in user density affects a group of cells (sky blue in the figure) around the target cell, including the comparing cell nearby. As the correlation level with the cell nearby is still preserved a false alarm could be avoided if this cell is chosen as comparing cell as no suspicion is raised by it. However, the correlation level with the cell far away is decreasing as the cell far away does not experience the load changes by the local peak and so a suspicion is raised also there. Hence, if the comparing cell had been chosen far away, then both sides raise a suspicion and we have a false alarm. In this case it is an advantage to choose the comparing cell nearby.

So, we can say, if we have to pick a comparing cell out of two possibilities, one nearby and the other far away, then there is no reason to choose the one far away.

In order to follow the second principle we consider situations where traditional monitoring methods can lead to false alarms, i.e. situations where a suspicion for a fault is raised but the target cell is not degrading. For instance, KPIs like load levels, number of unsuccessful call attempts, channel quality, etc. are monitored and traditionally suspicions are raised if they cross certain thresholds. Such situations can occur by unusual and hence unexpected user behavior. If for instance a peak in user density affects a group of cells it can be that the traditional performance indicators of some of the cells interpret this as being a cell performance degradation. However, the load correlation level might still be preserved as the unexpected peak in user density affects the whole group of cells. In that case the performance indicators raise a suspicion while the correlation level monitoring does not. Hence a combined logic would still not raise an alarm if the combined logic requires a suspicion from both sides. A false alarm would be avoided.

So, while the second principle suggests to choose comparing cells nearby and the first principle to choose them with correlation levels above a certain critical correlation level, we summarize the steps to choose comparing cells for the target cells:

1. Determine the critical load correlation level that comparing cells shall satisfy
2. For each target cell determine how many comparing cells shall be picked. The operator should decide the number of comparing cells. We recommend that picking two or three cells are enough to detect a degraded cell.
3. For each target cell monitor the load correlation levels with the cells in the first tier of surrounding cells
4. Check the consistency of the correlation between the target cell and the cells in the first tier. If the correlation between target cell and other cells remains higher than the critical correlation level then this cell shall be picked as comparing cell.
5. If there are not yet sufficiently many comparing cells picked so far go to the next tier and continue there, etc.

Steps 3 to 5 in our procedure therefore aim at choosing possible candidates near the target cell. If we find candidates we can stop, otherwise we enlarge the neighborhood to look for further cells. In this way comparing cells can be picked such that the two principles for the choice are followed as good as possible.

5 Conclusion and Future Work

In this paper we proposed a new method of cell degradation detection using the correlation between cells. In particular, the use of correlation between cell pairs is studied in detail and we used performance data collected from a real operational network in a European city for evaluation of the method. Our analysis suggests that the correlation coefficient between cell pairs can be utilized as a means for the detection of degradations in cells. The future work includes observing data in a real network together with the operators knowledge if problems in cells have occurred or not and to test the algorithm in this way.

References

[1] 3GPP TS 32.521, Self-Optimization OAM; Concepts and Requirements. Release 9, (June 2009)

[2] Hämäläinen, S., Sanneck, H., Sartori, S.: LTE Self-Organising Networks (SON): Network Management Automation for Operational Efficiency (January 2012)

[3] Ramiro, J., Hamied, K.: Self-Organizing Networks (SON): Self-Planning. Self-Optimization and Self-Healing for GSM, UMTS and LTE (January 2012)

[4] Cheung, B., Kumar, G.N., Rao, S.: Statistical Algorithms in Fault Detection and Prediction: Toward a Healthier Network. Bell Labs Technical Journal 9(4), 171–185 (2005)

[5] Rao, S.: Operational Fault Detection in Cellular Wireless Base-Stations. IEEE Transactions on Network and Service Management 3(2) (2006)

[6] Zanier, P., Guerzoni, R., Soldani: Detection of Interference, Dominance and Coverage Problems in WCDMA Networks. In: PIMRC (2006)

[7] Barreto, G.A., Mota, J.C.M., Souza, L.G.M., Frota, R.A., Aguayo, L., Yamamoto, J.S., Macedo, P.E.O.: Competitive Neural Networks for Fault Detection and Diagnosis in 3G Cellular Systems. In: Telecommunication and Networking –ICT (2004)

[8] Mueller, C.M., Kaschub, M., Blankenhorn, C., Wanke, S.: A Cell Outage detection Algorithm Using Neighbor Cell List Reports. In: IWSOS Proceedings of the 3rd International Workshop on Self-Organiznig Systems (2008)

[9] Turkka, J., Chernogorov, F., Brigatti, K., Ristaniemi, T., Lempiäinen, J.: An Approach for Network Outage Detection from Drive-Testing Databases. Journal of Computer Networks and Communications Article ID 163184, 13 pages (2012), doi:10.1155/2012/163184

[10] Asghar, M.Z., Hämäläinen, S., Meinke, N.: Experimental System for Self-Optimization of LTE Networks. In: 15th ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems (2012)

[11] Asghar, M.Z., Hämäläinen, S., Ristaniemi, T.: Self-Healing Framework for LTE Networks. In: 17th IEEE International Workshop on Computer-Aided Modeling Analysis and Design of Communication Links and Networks (2012)