A Cell Outage Detection Algorithm using Neighbor Cell List Reports

Christian M. Mueller, Matthias Kaschub, Christian Blankenhorn, and Stephan Wanke

Universität Stuttgart
Institute of Communication Networks and Computer Engineering
Pfaffenwaldring 47, 70569 Stuttgart, Germany
`[firstname.lastname]@ikr.uni-stuttgart.de, stephan.wanke@web.de`

Abstract. Base stations experiencing hardware or software failures have negative impact on network performance and customer satisfaction. The timely detection of such so-called outage or sleeping cells can be a difficult and costly task, depending on the type of the error. As a first step towards self-healing capabilities of mobile communication networks, operators have formulated a need for an automated cell outage detection. This paper presents and evaluates a novel cell outage detection algorithm, which is based on the neighbor cell list reporting of mobile terminals. Using statistical classification techniques as well as a manually designed heuristic, the algorithm is able to detect most of the outage situations in our simulations.

1 Introduction

1.1 Self-Organizing Networks

The configuration, operation and maintenance of mobile communication networks becomes increasingly complex, due to a number of different reasons. First, the continuous demand for bandwidth and coverage, driven by more and more widespread mobile Internet usage and new applications, urges operators to upgrade their backhaul networks and invest in new backhaul and air interface technologies. These need to be integrated into an existing infrastructure, which increases heterogeneity and in-turn complicates the management of these networks. Second, more sophisticated algorithms are implemented at the radio layers to exploit various physical effects of a wireless communications channel in order to increase system capacity. Third, more and more optimizations are conducted at cell-individual level to better exploit the particularities of a site. Given that nowadays, UMTS networks provide around 10^5 operator-configurable parameters [1], manual optimization already is a tedious and costly task.

Facing the challenge of increasingly complex network management, operators are at the same time under a tremendous cost pressure, driven by a strong competition among existing players and new entrants to the mobile communications sector. In this situation, the ability to efficiently manage their internal resources and the ability to cut down capital and operational expenditure becomes a competitive edge.

These constraints currently foster a trend towards a higher degree of automation in mobile communication networks. Under the term *Self-Organizing Networks (SON)*, the Next Generation Mobile Networks (NGMN) alliance has published a set of use cases for the application of self-configuration, self-optimization and self-healing concepts in deployment, planning and maintenance of mobile communication networks [2, 3]. The 3GPP has also recognized the need for more automation and cognitive functions in future wireless networks and has therefore started to work on concepts, requirements and solutions [4, 5] In addition, research projects are underway to develop solutions to these requirements [6, 7].

1.2 SON Use Case: Cell Outage Detection

One of the SON use cases listed in [3] concerns the automated detection of cells in an outage condition, i.e. cells which are not operating properly. The reasons for the erroneous behavior of an outage cell, often also denoted as sleeping cell, are manifold. Possible failure cases include hardware and software failures, external failures such as power supply or network connectivity, or even misconfiguration. Consequently, the corresponding observable failure indications are diverse, too. Indications can be abnormal handover rates, forced call terminations, decreased cell capacity or atypical cell load variations. While some cell outage cases are easy to detect by Operations Support System (OSS) functions of the network management, some may not be detected for hours or even days. Detecting those outage cases today is usually triggered by customer complaints. Subsequently, discovery and identification of the error involves manual analysis and quite often requires unplanned site visits, which makes cell outage detection a costly task.

A rough classification of sleeping cells is provided in [8]. According to this, a *degraded* cell still carries some traffic, but not as it would if it was fully operational. A *crippled* cell in contrast is characterized by a severely decreased capacity due to a significant failure of a base station component. Finally, a *catatonic* cell does not carry any traffic and is inoperable.

In our work, we focus on the detection of a catatonic cell. More specifically, our failure assumption is that the base station experiences a failure in one of its high-frequency components. This might be an error of the HF amplifier, antenna or cabling. The consequence is, that the cell is not visible anymore to users or neighbor base stations, i.e. it does not transmit a pilot signal anymore. However, from a network point of view, the cell appears to be empty, but still operational. As requirements to the detection algorithm, we have decided to only use already available measurement data and to avoid the need for new sensor equipment or the introduction of dedicated measurement procedures. Furthermore, the algorithm shall be able to reduce detection time to seconds or minutes.

While the basic idea and some preliminary results of our outage detection technique have already been presented in [9], here a a detailed description of the algorithm and an in-depth evaluation of its performance is provided. The remainder of this paper is structured as follows: Section 2 presents the outage detection algorithm. The evaluation methodology and system model are described in section 3. Section 4 examines the performance of the detection algorithm and presents simulation results. Section 5 finally summarizes and draws conclusions.

2 Cell Outage Detection Algorithm

The idea behind our outage detection algorithm is to use Neighbor Cell List (NCL) reports to create a graph of visibility relations, where vertices of the graph represent cells or sectors. Edges of the graph are generated according to the received NCL reports, where an edge weight determines the number of mobile terminals that have reported a certain neighbor relation. More details on how the graph is constructed will be given in section 2.1.

Neighbor Cell List reports are always generated during active connections, when mobile terminals continuously measure the signal strength and quality of the radio channel of the cell it is currently connected to. They also measure the signals of several neighbor cells, which constitute candidate cells for potential handovers. These measurements are transmitted to the current cell to decide whether a handover might have to be performed. In GSM, for example, a terminal in an ongoing connection measures up to 16 base stations within its visibility range and sends NCL reports at 480 ms intervals, containing the best 6 measurement results [10]. The NCL reports of different terminals are retrieved by a subordinated network entity, e.g. a radio network controller in UMTS, respectively a base station controller in GSM. While measurement reporting details differ from GSM to UMTS or LTE, we only assume that lists of neighbor cells are reported at regular intervals, which can be realized in all the different radio access technologies.

2.1 Monitoring the Visibility Graph

Monitoring changes in the visibility graph is a key element of our detection algorithm. The visibility graph is created at regular intervals and any unusual variations of the graph might indicate a cell outage. Figure 1 depicts examples of such a graph. Comparing two successive graphs $G(t_1)$ and $G(t_1 + T)$, the detection algorithm is sensitive towards any nodes becoming isolated in $G(t_1+T)$, in which case a so-called *change pattern* is created.

In Fig. 1(a), a mobile terminal UE_A is currently being served by sector 2, while UE_B is in an active connection currently being handled by sector 3. The visibility range of both mobiles in this simplified example is denoted as a circle around their respective positions. Thus, the NCL report of UE_A contains cell 2 as the currently serving cell and cells 1 and 3, which are neighbor cells within its visibility range. The NCL report of UE_B is constructed accordingly. In the resulting visibility graph, the NCL report of UE_A now leads to the two directed edges (blue) originating from node 2 to the nodes 1 and 3. Analogously, the NCL report of UE_B results into the two edges (red) originating from node 3.

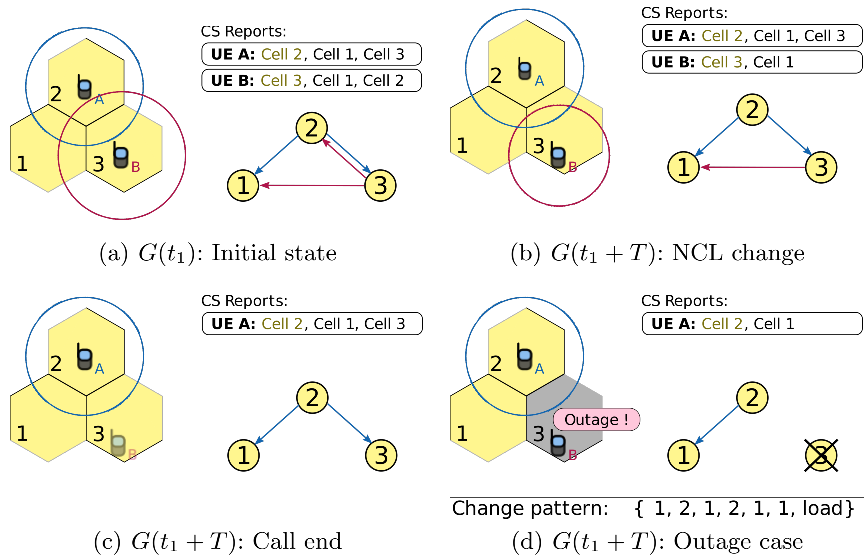

(a) $G(t_1)$: Initial state (b) $G(t_1 + T)$: NCL change

(c) $G(t_1 + T)$: Call end (d) $G(t_1 + T)$: Outage case

Fig. 1. Visibility Graph examples with two terminals being served by cells 2, respectively 3

Every edge in the graph has a certain weight, representing the number of UEs which have reported a neighbor relation between a pair of cells. In the example in Fig. 1, all edge weights equal one and are therefore not shown.

During normal operation of the network, the visibility graph is subject to frequent changes, caused by starting and ending calls, user mobility, changes in radio propagation and changes in the NCL reporting itself. Figures 1(b) and 1(c) show typical examples of graph variations. In Fig. 1(b), UE_B has changed its position such that cell 2 is no longer in its visibility range. The corresponding NCL reports are depicted on the right hand side of Fig. 1(b), together with the resulting topology of the visibility graph. Another variation of the graph is depicted in Fig. 1(c), where UE_B has terminated its connection and does not send NCL reports anymore. Consequently, the visibility graph now only contains the neighbor relations reported by UE_A.

Figure 1(d) finally depicts the resulting visibility graph if cell 3 experiences a failure and becomes catatonic. The NCL reports of UE_B are not received anymore. At the same time, cell 3 disappears from the candidate set of UE_A. The resulting graph now only contains a single edge from node 2 to node 1. Note that every outage situation results into an isolated node in the visibility graph.

Figure 1(d) shows the change pattern of the graph variation compared to the initial state of the graph in Fig. 1(a). A change pattern is a 7-tuple with the following attributes, that is created for every isolated node:

– number of disappeared edges to the isolated node
– number of disappeared edges from the isolated node

– sum of the edge weights to the isolated node
– sum of the edge weights from the isolated node
– number of UEs of which no reports are received anymore
– number of UEs with changed neighbor cell list reports
– load as sum of the number of active calls of the direct neighbor cells

The respective elements of the 7-tuple are the features or input parameters of the classification step described in the following section 2.2.

The change patterns that are created when nodes become isolated are not necessarily unique. In the sample scenario in Fig. 1, the same pattern would have been created if a movement of UE_A and the end of a connection of UE_B occurred within a single monitoring interval. Thus, an isolated node pattern in the visibility graph is a necessary condition for an outage situation, but not a sufficient one. However, our assumption is that cell outages create characteristic change patterns which, in many cases, can be distinguished from normal fluctuations of the visibility graph. The outage detection problem thus translates into a classification problem on change patterns of the visibility graph.

2.2 Outage Detection as a Binary Classification Problem

The classification of a set of items is the assignment of similar items to one out of several distinct categories. The task of separating change patterns into outage and non-outage situations can be regarded as a binary classification problem with a set of predefined classes.

Classification algorithms are widely used in the field of automatic pattern recognition [11]. A classifier can either be knowledge-based (e.g., expert systems, neural networks, genetic algorithms) or make use of mathematical techniques from multivariate statistics (e.g. cluster analysis, classification and regression trees). We applied three different classification techniques, a manually designed expert system and two others using statistical classification techniques. In statistical classification, the n-dimensional tuple space is separated according to statistical measures, which allows to automatically construct a classifier from a given training set of patterns. A large number of statistical methods exist to infer the classification rules, of which an iterative decision-tree algorithm and a linear discriminant analysis have been applied here and will be further detailed in the following:

Expert System. From the class of knowledge-based classifiers, an expert system has been applied to the outage detection problem. A set of rules has been manually constructed, based on the observation of the change patterns of the visibility graph. The rules are expressed as a sequence of *if ... then* statements, with threshold values tailored to our evaluation scenario. The expert system has been designed such that it performs a conservative classification, i.e., the number of false alarms shall be minimized.

Decision-Tree (DT). A tree of binary decisions on the attributes of the respective patterns is automatically created from a training set. The decision about when to split into separate branches is thereby determined by the Gini impurity function, which is used by the classification and regression tree (CART) algorithm [12]. Similar to the expert system, the outcome of the DT algorithm can be regarded as a sequence of *if . . . then* statements, with the difference that now these rules are being determined automatically.

Linear-Discriminant Function (LDF). The LDF classifier belongs to the class of linear classification algorithms. The discriminant function is a linear combination of the pattern's attributes, whose coefficients are determined from the training set using a least square estimation [13].

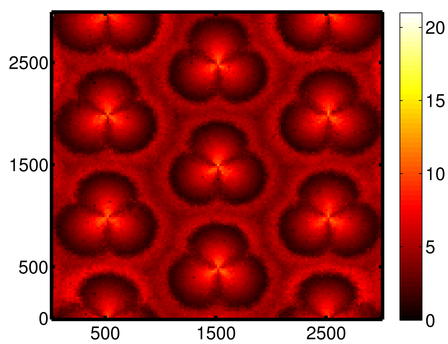

Fig. 2. SINR without outage

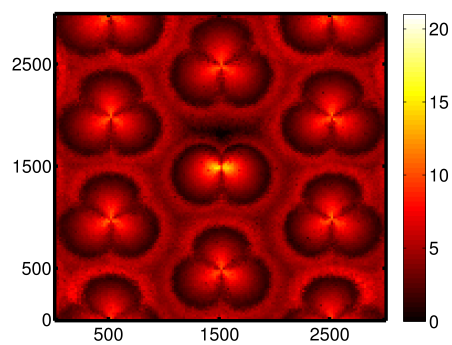

Fig. 3. SINR in outage case

3 Evaluation Methodology

The performance evaluation of the cell outage detection algorithm involves several steps. First, a system-level simulation of a multiple sites scenario is run several times for a duration of 100.000 s without any outage situation, in order to collect change patterns from the normal operation of the network. Second, a Monte Carlo-like simulation is performed to collect change patterns for cases where a cell becomes catatonic. For each of the 3000 drops of this Monte Carlo-like simulation, mobile terminals were positioned randomly and the change patterns are obtained immediately after a cell sector is configured to be in an outage state. Finally, a training set and a test set is constructed from the collected patterns and fed to the classification algorithms. For the statistical classification, the *LSS Classification Toolbox* in its default parametrization is used [14]. The system-level simulations have been conducted with the event-driven IKR Simlib simulation library [15].

The scenario used in the event-driven simulations consists of 19 sites, respectively 57 cells, with a base station distance of 1000 m. The system was modeled as an FDM system with a maximum capacity of 20 simultaneously active calls per cell. A fractional frequency reuse is assumed with half the capacity reserved for the re-use 1 region in the cell center and the remaining capacity distributed over the re-use 3 regions at the cell border. Inter-cell interference is considered as proportional to the actual load in the corresponding re-use areas of neighbor cells. The Walfish-Ikegami model (non-LOS) determines the path loss and spatially correlated shadowing planes similar to [16] are used. A global shadowing plane models the influence of buildings and other common obstructions to radio propagation, while local shadowing planes for each base station account for the different antenna locations. The 3-sector antenna pattern corresponds to [17]. A handover algorithm is implemented which is sensitive towards SINR and RSSI (Received Signal Strength Indicator) values. The measurement reporting of neighbor cell lists is assumed to be GSM-like with a reporting period of 480 ms, which in our setting corresponds to the graph update period. Further parameters are given in Table 1. SINR plots extracted from the simulation for an outage and a non-outage case for an average cell load of 25% are given in Figures 2 and 3, respectively.

Table 1. Scenario and simulation parameters

Parameter	**Value**
Base station distance	1000 m
Shadowing std. dev. (global)	7.0 dB
Shadowing std. dev. (local)	1.0 dB
Shadowing correlation length	50 m
BS power to individual UE	35 dBm
Signal detection threshold at UE	-96 dBm
Traffic model (Poisson arrivals)	$\lambda = \frac{1}{100\,\mathrm{s}}$
Call holding time (neg.-exp.)	$h = 35\,\mathrm{s}$
User movement	Random Direction
User speed	$v = 1 - 5\,\frac{\mathrm{m}}{\mathrm{s}}$
Segment length (uniform)	$l = 20 - 800\,\mathrm{m}$
Direction angle (uniform)	$0 - 360$

4 Analysis and Results

4.1 Outage Observability

The quality of the outage detection is largely determined by the performance of the classification algorithm. However, even with perfect classification, some outage cases might still not be detected. In particular in low load situations when user density is scarce, a cell might not serve any users and might not be measured by any of the UEs of the surrounding cells. A cell outages thus does not result into changes of the visibility graph. In this case, the outage would not be detectable at all.

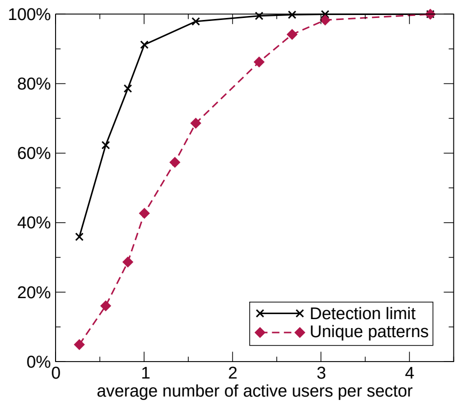

Fig. 4. Upper bound of cell outage observability and ratio of unique patterns

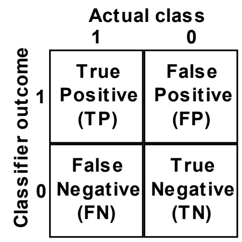

Fig. 5. Confusion Matrix

Figure 4 depicts the upper bound of detectable outages for the evaluation scenario. The black solid line denotes the ratio of detectable outages over the average number of active connections per cell. It can be seen that in low load conditions, a large proportion of the cell outages have no representation in the visibility graph. Starting from an average 1.6 active connections per sector, which in our setup corresponds to a cell load of 10%, already more than 95% of the outages can theoretically be detected. As outlined in section 2.1, outages do not necessarily create unique change patterns because normal fluctuations of the visibility graph might lead to the exact same patterns. The dashed red curve denotes the proportion of patterns uniquely identifying a cell outage. This gives a reference value which constitutes the performance achievable by a hypothetical look-up table based detection algorithm, assuming that all possible outage patterns are known in advance.

4.2 Classification Quality

The quality of a classification algorithm is characterized by the so-called confusion matrix (see Fig. 5).

True Positive. A *TP* denotes the case that an outage (i.e. a cell becomes catatonic) has occurred and it has been successfully detected.

False Negative. A *FN* denotes the case that an outage has occurred, but it has not been detected.

False Positive. A *FP* occurs, when there is no catatonic cell but the detection algorithm nevertheless reports an outage.

True Negative. A *TN* is when there is no outage and the algorithm correctly recognizes the change pattern as normal fluctuation of the visibility graph.

TP, *FN*, *FP*, and *TN* denote the total numbers of occurrences for these four cases in a simulation run. From the confusion matrix, a number of metrics can be derived. On the one hand, the *Sensitivity* or *True-Positive-Rate* R_{TP} denotes whether a classifier is able to correctly detect the outage patterns from the set of all patterns:

$$R_{\mathrm{TP}} = \frac{TP}{TP + FN} \tag{1}$$

On the other hand, the *Fall-out* or *False-Positive-Rate* R_{FP} gives the tendency of a classification algorithm to create false alarms:

$$R_{\mathrm{FP}} = \frac{FP}{FP + TN} \tag{2}$$

Achieving a high sensitivity generally comes with an increase in the rate of false alarms. The two metrics can therefore be regarded as complementary. A so-called *Receiver-Operating-Characteristics* (ROC) plot is a way to visualize the performance of binary classification algorithms regarding these two metrics and is primarily used in signal detection theory [18]. In a ROC plot, the sensitivity is drawn over the false-positive-rate. A perfect classification would thus result in a mark in the upper left corner of the plot. The line through the origin gives the reference values for a random decision, where each of the classes is chosen with probability one half.

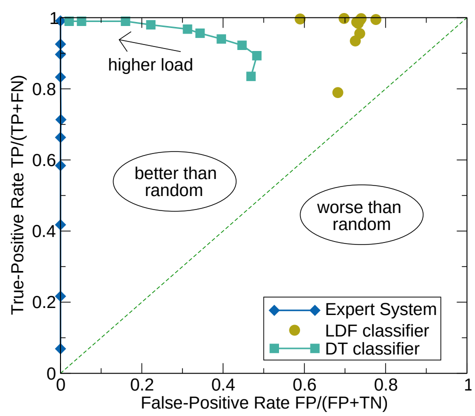

Fig. 6. ROC plot of the classification algorithms

Figure 6 shows a ROC plot of the classification algorithms applied here, for the range of cell load values also used in the other figures. After having been trained with a training set, the classifiers had to determine the correct mapping for test sets, where each of the patterns is tested only once. In doing so, Fig. 6 is independent of the frequency of occurrence of the different patterns in the outage detection scenario. It thus solely gives the performance of the classification algorithms, which is not identical to the resulting performance of the outage detection algorithm. It can be observed that the Expert System and the Decision Tree classifiers are able to achieve a high sensitivity, especially for higher load scenarios, which is consistent with the results from Fig. 4. The Expert

System has been designed to perform a conservative detection. In other words, the Expert System is more likely to indicate False Negatives than False Positives. Compared to this, the DT classifier is characterized by a much higher risk to create false alarms. Finally, the LDF classifier has high false alarm probability and, in contrast to the other algorithms, does not show any improvement with increasing system load.

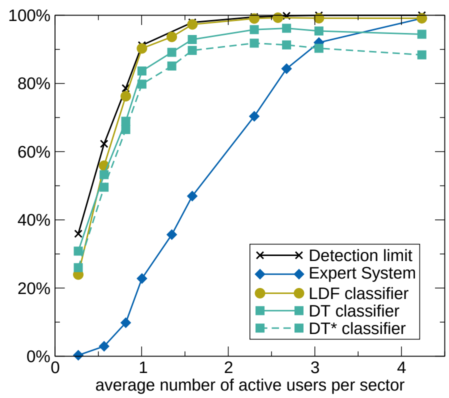

Fig. 7. Sensitivity of the classification algorithms

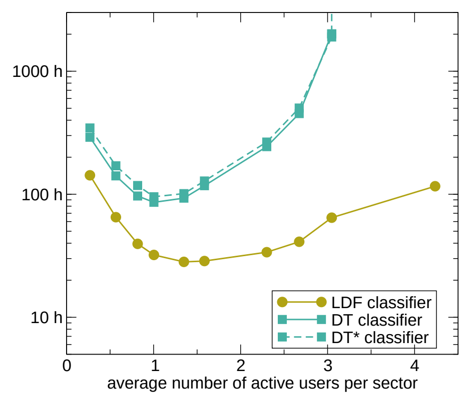

Fig. 8. Mean time between false alarms per cell sector

4.3 Outage Detection Performance

In order to determine the performance of the outage detection algorithm, not only the classification quality, but also the problem-specific frequency of occurrence of the patterns has to be taken into account. A wrong classification of frequently occurring patterns could significantly deteriorate the suitability of the classification algorithms to the outage detection problem, whereas a classification error for rare patterns does not have much influence. Figure 7 depicts the sensitivity of the classification algorithm with respect to the observability bound presented in Fig. 4. It thus gives a measure of the proportion of outages that can be detected with the approach presented here. Using the statistical classifiers, sensitivity is close to the maximum achievable sensitivity. The Expert System only attains moderate sensitivity due to its conservative decision rules.

Finally, Fig. 8 depicts the fall-out risk per cell in terms of the mean time between false alarms. Only the statistical classifiers are shown, given that the Expert System did not produce any false positive classifications over the whole monitoring period. For lower load, overall sensitivity is low and therefore the risk of rising false alarms also is lower than in the middle range of the curves in Fig. 8. However, a mean time between false alarms of only several tens or hundreds of hours would already result into several false alarms per cell being raised within a single week. Even though the applied classification algorithms here provide means to mitigate this problem by giving higher penalty to false positive classification, the results do not differ significantly. An example is shown in Figures 7 and 8 for a modified decision tree classifier (denoted as DT$*$), which has given a ten times higher penalty for false positive classifications.

4.4 Discussion

From the previous subsections, it can be summarized that although the statistical classifiers show an overall good detection sensitivity, tendency towards generation of false alarms is still far too high for practical applications. This is mainly due to the partly ambiguous change patterns and the non-linearity of the problem. The expert system achieves much better performance here, but generally suffers from the need to manually tune its threshold parameters to the scenario under investigation. These observations motivate further work, employing more flexible classification algorithms such as fuzzy classifiers or Support Vector Machines [19].

From a network operation point of view, a detection algorithm as described here can be applied to whole or parts of a network. That is, collection of NCL reports and its interpretation can be done in an RNC or BSC, without the need for additional signaling. For LTE, the scope of a single eNodeB is too small and a master eNodeB or another management device would have to monitor the visibility graph. Although this results in additional signaling, the NCL reports of each terminal in a cell can be aggregated by the eNodeB over one monitoring period, which results in one signaling message per monitoring interval being sent to the respective management device.

5 Conclusion

Automated detection of base station failures or sleeping cells is a prerequisite for future self-healing capabilities in mobile communication networks. The detection of some failure conditions still imposes a challenge to network operators and demands for additional means to observe base station behavior from the outside. An algorithmic detection without the need for additional sensor equipment thereby represents the most cost effective way of improving outage detection capabilities in existing networks.

The presented approach of using neighbor cell measurement reports of mobile terminals to determine outage situation showed good performance, even in moderate cell load conditions. However, the relatively high risk of false alarms still prevents from practical application. Although the algorithm is able to quickly detect a large proportion of the outages in a certain cell load range, it has been shown that there is no possibility to detect outages when no active users are close to the respective cell. In this case, the presented algorithm may be combined with other detection mechanisms operating on a longer time scale.

References

1. Lehser, F.: Self Organising LTE/SAE Network: Operator Requirements and Examples. 20th workshop of the ITG group 5.2.4, Ulm, Germany, available online at http://www.ikr.uni-stuttgart.de/Content/itg/fg524/ (September 2006)
2. NGMN Alliance: Use cases related to self organising network: Overall description. www.ngmn.org (April 2007)
3. NGMN Alliance: Annex A (informative) of use cases related to self organising network: Overall description. www.ngmn.org (April 2007)
4. 3GPP: Telecommunication management; Self-Organizing Networks (SON); Concepts and requirements. TS 32.500 (2008)
5. 3GPP: Evolved Universal Terrestrial Radio Access Network (E-UTRAN); Self-configuring and self-optimizing network (SON) use cases and solutions. TR 36.902 (2008)
6. FP7 Socrates: Self-optimisation & self-configuration in wireless networks. http://www.fp7-socrates.org/
7. FP7 E3: End-to-end efficiency (e3). http://ict-e3.eu
8. Cheung, B., Fishkin, S.G., Kumar, G.N., Rao, S.: Method of monitoring wireless network performance. USPTO Patent Application 20060063521 (March 2006)
9. Mueller, C.M., Kaschub, M., Blankenhorn, C., Wanke, S.: Design and evaluation of detection algorithms for base stations outages. In: 11th COST 290 MCM, TD(08)003, Tampere, Finland (2008)
10. Eberspächer, J., Vögel, H.J., Bettstetter, C.: GSM Global System for Mobile Communication. B.G. Teubner Stuttgart, Leipzig, Wiesbaden (2001)
11. Jain, A., Duin, R., Mao, J.: Statistical pattern recognition: a review. IEEE Transactions on Pattern Analysis and Machine Intelligence **22**(1) (2000) 4–37
12. Breiman, L., Friedman, J., Stone, C.J., Olshen, R.: Classification and Regression Trees. Chapman & Hall/CRC; 1 edition (1984)
13. Duda, R.O., Hartr, P.E.: Pattern Classification and Scene Analysis. John Wiley and Sons Ltd. (1973)
14. Uhlich, S.: Classification Toolbox V1.2. available online at http://www.lss.uni-stuttgart.de (March 2008)
15. IKR: Simulation library v2.6. available online at: http://www.ikr.uni-stuttgart.de
16. Forkel, I., Schinnenburg, M., Ang, M.: Generation of two-dimensional correlated shadowing for mobile radio network simulation. In: Proceedings of the Wireless Personal Multimedia Communications (WPMC). (2004)
17. NGMN Alliance: Radio access performance evaluation methodology. www.ngmn.org (June 2007)
18. Fawcett, T.: ROC Graphs: Notes and Practical Considerations for Researchers. Technical report, Palo Alto, USA: HP Laboratories (2004)
19. Wang, L., ed.: Support Vector Machines: Theory and Applications. Studies in Fuzziness and Soft Computing. Springer Berlin (2005)