Improving Cell Outage Management through data analysis

I. de-la-Bandera*, P. Muñoz*, I. Serrano† and R. Barco*

*Universidad de Málaga, Departamento de Ingeniera de Comunicaciones.

29071. Málaga, Spain.

Email: {ibc, pabloml, rbm}@ic.uma.es

† Ericsson, SDT EDOS-DP. 29590 Campanillas, Málaga, Spain.

Email: inmaculada.serrano@ericsson.com

Abstract

Self-Organizing Networks (SON) is an important feature for network management automation in the new generation of mobile communications. While SON have been considered as part of the recent 3rd Generation Partnership Project (3GPP) standards such as Long Term Evolution (LTE), it is expected that next 5th Generation (5G) mobile networks present new challenges for SON solutions. One of the most important use cases in Self-Healing is Cell Outage Compensation (COC). This paper proposes an important improvement of COC function based on analyzing large real data sets from live networks in order to adapt compensation to the real neighboring context. First, two methods (offline and online) to classify outages depending on the degraded metrics in neighboring cells are proposed and results for a set of cell outages occurred in a live LTE network are presented. Second, a method for estimating the lost traffic due to cell outages is proposed in order to quantify the load that cannot be intrinsically absorbed by neighboring cells. Finally, a novel COC methodology is proposed by taking into account the results obtained in the two previous studies.

Keywords: LTE, Self-Healing, Cell Outage Compensation, Self-Organizing Networks.

1 Introduction

The complexity of network operation and cost savings are foreseen to be the biggest challenges in 5th Generation (5G) of mobile radio access technology which can be tackled by two recent technologies: Self-Organizing Networks (SON) and Big Data. On the one hand, SON is a well-known concept for network management automation which has been standardized by the 3rd Generation Partnership Project (3GPP) and extensively researched for newly deployed technologies such as Long Term Evolution (LTE) [1]. On the other hand, Big Data is an evolving term that can be applied in this context to describe large amounts of heterogeneous data collected from User Equipments (UEs) and network elements which can be used to optimize management tasks, increasing automation and reducing costs [2].

Currently, LTE is increasingly becoming a mature technology, as LTE networks are already being massively deployed around the world by most service providers. This opens the doors for developing smart algorithms for any of the categories defined in SON: Self-Configuration, Self-Optimization and Self-Healing. Traditionally, a significant effort has been devoted to the investigation of SON solutions based on analytical and simulated models. However, the main challenges are now linked to the exploitation of the available data in order to create a useable knowledge base that maximizes the performance of current SON solutions. This is particularly interesting in the area of Self-Healing [3] which comprises the SON functionalities in charge of failure detection, diagnosis, compensation and recovery with the aim of managing major service outages and degradations. In this context, the availability of historical data is crucial to understand how the network behaves under normal and faulty conditions or how the faults impact on the quality-of-experience.

Within Self-Healing, the use cases mostly addressed in the literature are the Cell Outage Detection (COD) and the Cell Outage Compensation (COC), that have been specified by the 3GPP [4]. The importance of the cell outage fault relies on its high probability of occurrence and its critical impact on neighbor cells. Although they have been widely studied in the literature [5][6][7][8], the potential of data analysis techniques can provide further improvements to these functions. In the case of COD, the introduction of smart data analysis allows to predict the problem from its origins, and then take preemptive actions to avoid the occurrence of the problem [2].

In the case of COC, the analysis of massive amount of data can provide several benefits, which will be explored in this paper, together with new avenues and challenges for improving the COC by analyzing real data sets of cell outages. In particular, based on the analysis of data collected in neighboring cells, the main contributions of this paper are:

- First, the idea of addressing each neighboring cell of a cell in outage as an independent case, as opposed to the state-of-the-art approach of not considering the different effects of the outage on neighbors. It is based on the fact that cell outages can impact on each neighboring cell in a different way, e.g. in some cases a cell outage could lead to mobility issues in neighboring cells due to a lack of a dominant cell, whereas in other cases the cell outage may produce a coverage hole. The analysis is made by measuring the correlation (dependence) between network metrics, called Key Performance Indicators (KPIs), in order to extract valuable information from a vast amount of data. Depending on whether the data set includes historical (offline) or real-time (online) cases of cell outages, different methods to estimate the impact of outages on neighboring cells have been proposed. To illustrate these approach, an analysis of cell outages in a mature LTE network has been carried out by looking at degraded metrics in neighboring cells.

- Second, the idea that estimating the amount of traffic that is potentially lost as a consequence of cell outages should be considered by COC policies. This is important since lost traffic can result in a potential loss of market share and revenues for the service provider. A method to estimate the lost traffic based on computing the traffic absorbed by neighboring cells is proposed.

- Third, by taking advantage of the previous cell outage analysis, a new concept of COC from the perspective of the neighboring cells is discussed. Current COC solutions are typically focused on modifying a specific radio parameter (e.g. the antenna downtilt or the transmitted power) in a predetermined number of neighboring cells to solve in most cases a problem of coverage hole. In this paper, the shortcomings of this kind of solutions are addressed and the benefits and challenges of using the proposed method to improve the COC algorithms are discussed.

With regards to the two last contributions above, to the authors' knowledge, there are no other works that calculate the traffic lost due to a cell outage. In addition, no other works presented a COC methodology that is able to adapt the compensation to the specific degradation produced by the cell outage.

2 Analysis of cell outages

2.1 Motivation

Once cell outages are detected, this paper aims to distinguish these outages by analyzing the impact on neighboring cells. This is justified by the fact that some cases will require more attention than others from the operator perspective. For example, depending on the amount of traffic absorbed in the affected area, neighboring cells could become overloaded or not. A slight increase in the traffic of neighboring cells may happen in deployments with high cell densification or in areas with scarce offered traffic. In addition, in early deployment stages, a cell outage may produce important coverage holes which would negatively impact on user accessibility and retainability. Conversely, in mature deployments, the new scenario after a cell outage may lead to mobility problems in neighboring cells, especially between those with overlapped coverage areas. Moreover, the type of scenario may determine the effects caused by a cell outage so that it is important to consider the type of affected cells (e.g. if the cell outage occurs in a HetNet). COC algorithms presented in the literature are based on the modification of one or more configuration parameters. For instance, the authors of [7] propose a COC algorithm based on modifications of the transmission power of the base stations and reallocating the users from the outage area. The algorithm presented in [8] is based on modifying the antenna tilt and the transmission power simultaneously. In all cases, the proposed algorithm is supposed to be applied to every outage failure that occur in the network. However, the same failure (i.e. a cell outage) may produce different effects in the neighboring cells that should be compensated in a different way. Fig 1 shows the effects produced in two different neighboring cells by the same cell outage. Specifically, the figure shows KPIs related to traffic, bad coverage and load for the two selected neighboring cells. It can be seen that the cell outage produces a degradation in the coverage and an increase in traffic for one neighboing cell while the other suffers a congestion problem. In this situation, a certain COC strategy as antenna tilt modification can improve the bad coverage problem but it may not have effect in the nighboring cell with the congestion.

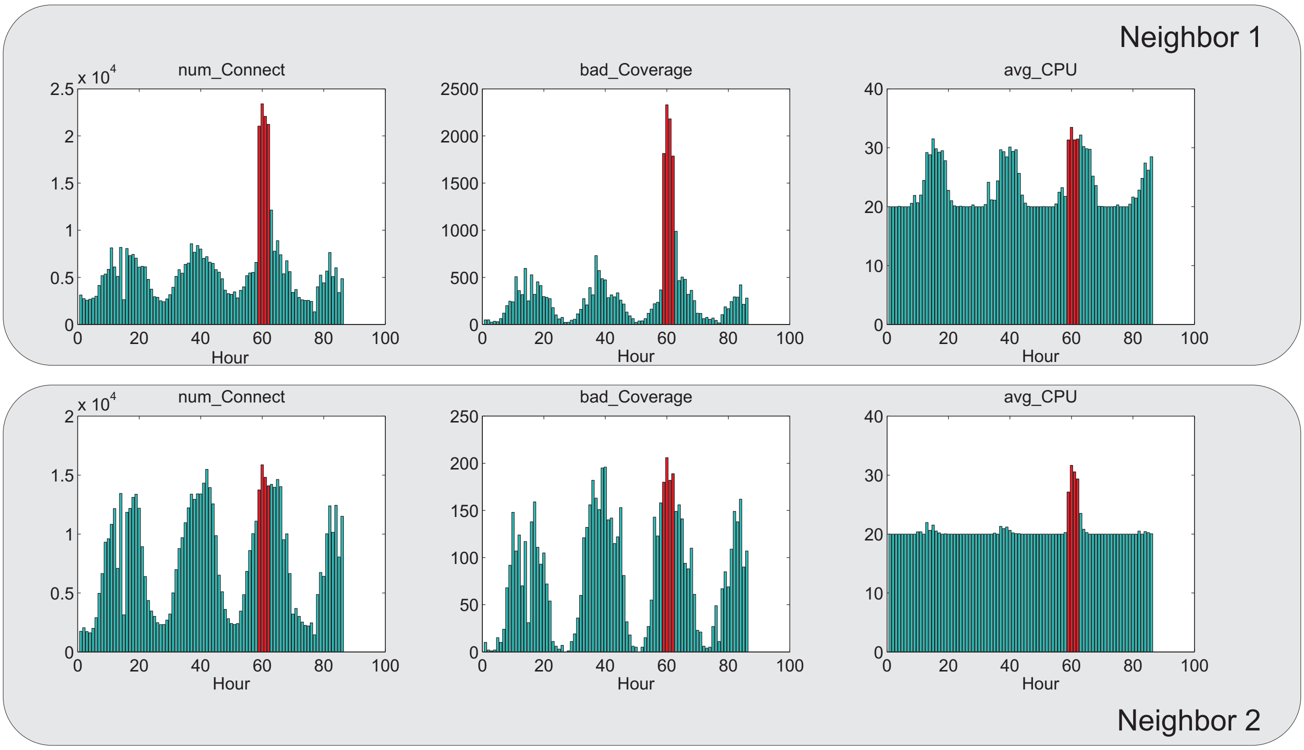

Figure 1: Degraded KPIs from neighboring cells in case of cell outage

In this section, the analysis of cell outages is based on identifying degraded KPIs in neighboring cells. In a first (offline) approach, a method for analyzing historical records of cell outages is proposed. Secondly, various methods to analyze cell outages in a real-time (online) SON fashion are presented and evaluated taking the offline approach as baseline.

2.2 The offline method for cell outages analysis

In Self-Healing, a common technique to manage anomalies in the network is to study the correlation between different metrics such as alarms or KPIs. In particular, several works using the correlation for COD algorithms can be found in the state-of-the-art [9][10]. The basis of these algorithms is the recognition of the characteristic behavior of faults by correlating the KPIs with each other. However, in this work, the correlation is not used to detect cell outages. Once the cell outage has been detected by a COD algorithm such as [5], the correlation is used to analyze the effects produced by the cell outage in the neighboring cells. Since the cell outage is a special fault in which KPIs from the problematic cell are lost, in that case the analysis of KPIs should necessarily be performed in neighboring cells.

The first step in correlation-based Self-Healing is the selection of metrics to be correlated. A representative set of KPIs to analyze degradation in cell outages is the following:

- *num_Connect* : it measures the number of established connections in a cell at the network layer, so that it is an estimation of the carried traffic.
- *bad_Coverage*: it is based on the number of UE reports indicating a low signal level received from the serving cell. A high value means that there is a lack of coverage in that cell.
- *inter_RAT_HO*: it calculates the number of calls that have been handed over (HO) to another cell belonging to a different radio access technology (RAT). This may reveal for example a lack of coverage in a certain RAT.
- *HO_PP*: it measures the number of HO ping-pong (*HO_PP*) events produced in the network. As a consequence of a wrong configuration of mobility parameters, an HO ping-pong event is given by two subsequent HOs between the source and the target evolved Node B (eNB) and vice versa. This problem may significantly decrease the performance of HOs.
- *avg_RSSI*: this KPI calculates the average value of the Received Signal Strength Indicator (RSSI), which includes energy not only from the desired signal but also from background noise and interference.

- *avg_CPU_Load*: it measures the average load carried by the Computer Processing Unit (CPU) in the eNB. This KPI may indicate hardware congestion.
- *num_Drops*: it counts the number of dropped calls in a cell.
- *Failed_Conn_Estab*: it measures the number of failed connection establishments as an indication of user accessibility problems.

The above KPIs are measured every hour by the Operations Support System (OSS). In the proposed method the time evolution of these metrics is used to find potential degradations on neighboring cells due to the impact of cell outages occurred in the network. Assuming the duration of analyzed cell outages to be only a few hours (e.g. less than 24 h), the procedure to detect KPI degradations is shown in Fig. 2. The objective of the proposed algorithm is to generate a set of reference signals for each analyzed KPI that includes a certain level of degradation independiently of the values of the real KPI. These estimated signals will be compared to the original KPI to determine if there are degradations due to the outage. The algorithm is divided in some phases:

1. Define Observed Signal: For each neighboring cell, a set of samples of each KPI in a period of two days is collected. One day shoud comprise the cell outage and the other day should not be affected.
2. Estimate signal with normal behaviour. The previous signal is modified in order to generate a signal with no effects due to the cell outage. To achieve this, the samples that fall within the cell outage interval are substituted by samples at the same hours of the other day. These samples should not be affected by the cell outage.
3. Scale signal: The samples at the hours of the cell outage are scaled in order to simulate a possible degradation produced by the outage. In particular, two different scaling factors (e.g. 3 and 30) are applied to cover both slight and strong degradations. The resulting signals are the Reference Signals.

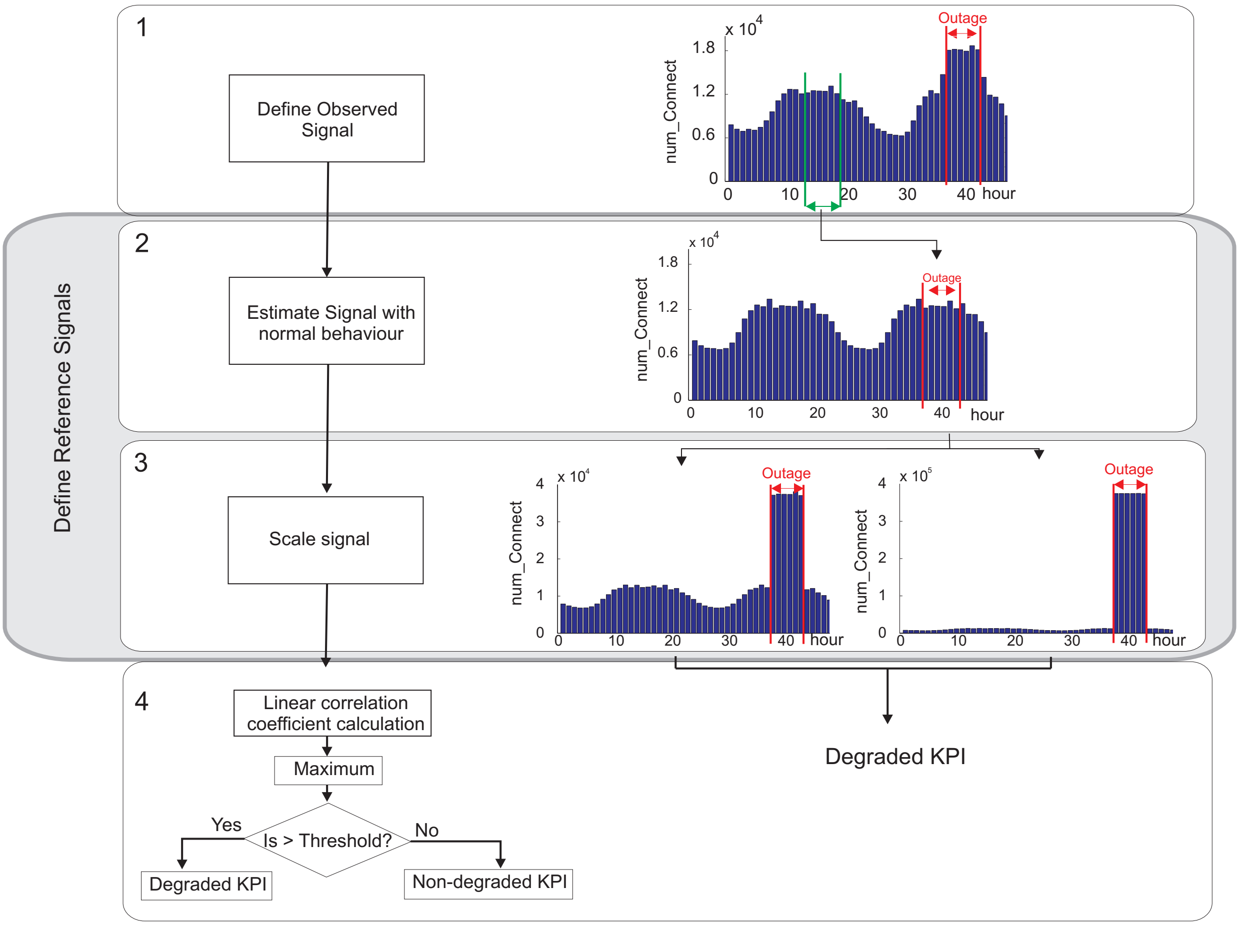

Figure 2: Diagram of the proposed method

4. Calculate correlation: The linear correlation coefficient is calculated between the observed and the reference signals. This is made for each used scaling factor. The maximum correlation value for the used scaling factors is selected. If the correlation value is above a specific threshold (e.g. 0.80), the KPI of the analyzed neighboring cell is considered to be degraded by the cell outage.

The time needed to complete the phases above is negligible comparing to the KPI updating periodicity so that this analysis does not increase significantly the total time of the compensation process.

In this algorithm, the time period of the signals has been carefully selected, since the correlation values may be very sensitive to this parameter. More specifically, the time length of the signals should be consistent with the duration of cell outages. For example, an excessive time length of the signals compared to the duration of cell outages could mask the impact of the cell outage on the correlation values. Another issue is related to the selection of days in the step 1 of the algorithm. Depending on the day of the week, the traffic pattern may be different, especially between workweek and weekend. Since hours at different days will be correlated, the selected days of the week should also be consistent in relation to the traffic pattern. To avoid such an issue, in this work, only time intervals during the workweek have been analyzed.

The proposed algorithm has been applied to 22 real cases of LTE macro cell outages. The selected dataset corresponds to an LTE network composed of 8000 cells approximately. The information about different KPIs is collected every hour. The inputs of the algorithm are the above mentioned KPIs which are measured in the neighboring cells of each outage cell. To consider a cell as neighbor, the number of HOs between both cells must be higher than zero. In addition, to simplify the analysis, only the three most degraded neighboring cells (i.e. those with higher number of degraded KPIs) are selected. In Fig. 3, a histogram of the different degraded patterns that have been found in neighboring cells is depicted. As observed, the most repeated pattern (i.e. pattern 1) is the one without any degradation in the KPIs. This reveals that in many cases cell outages only have an impact on none, one or two neighboring cells. Pattern 2 is given by the degradation in *num_Connect* and *bad_Coverage*, which are the typical effects of cell outages in traditional single-RAT networks, i.e. an increase of traffic in neighboring cells and the creation of a coverage hole. One of the most repeated patterns (4) is given by a degradation only in *HO_PP*. This may be common in networks with high degree of cell overlapping. In these cases, a cell outage may lead to a situation of a lack of dominant cell, where the number of unnecessary HOs would be increased. Pattern 5 presents a typical situation of multi-RAT networks, where the traffic can be mainly absorbed by other RATs (especially if cells are co-sited), impacting on *Inter_RAT_HO* rather than *num_Connect* of the neighboring cell. Patterns 3 and 6 are subcategories of pattern 2, which are typical effects of cell outages, as previously mentioned.

As the impact of cell outages on neighboring cell becomes more severe, the number of degraded KPIs is also higher. An example of this is represented by patterns 11, 13, 23, 24, 26 and

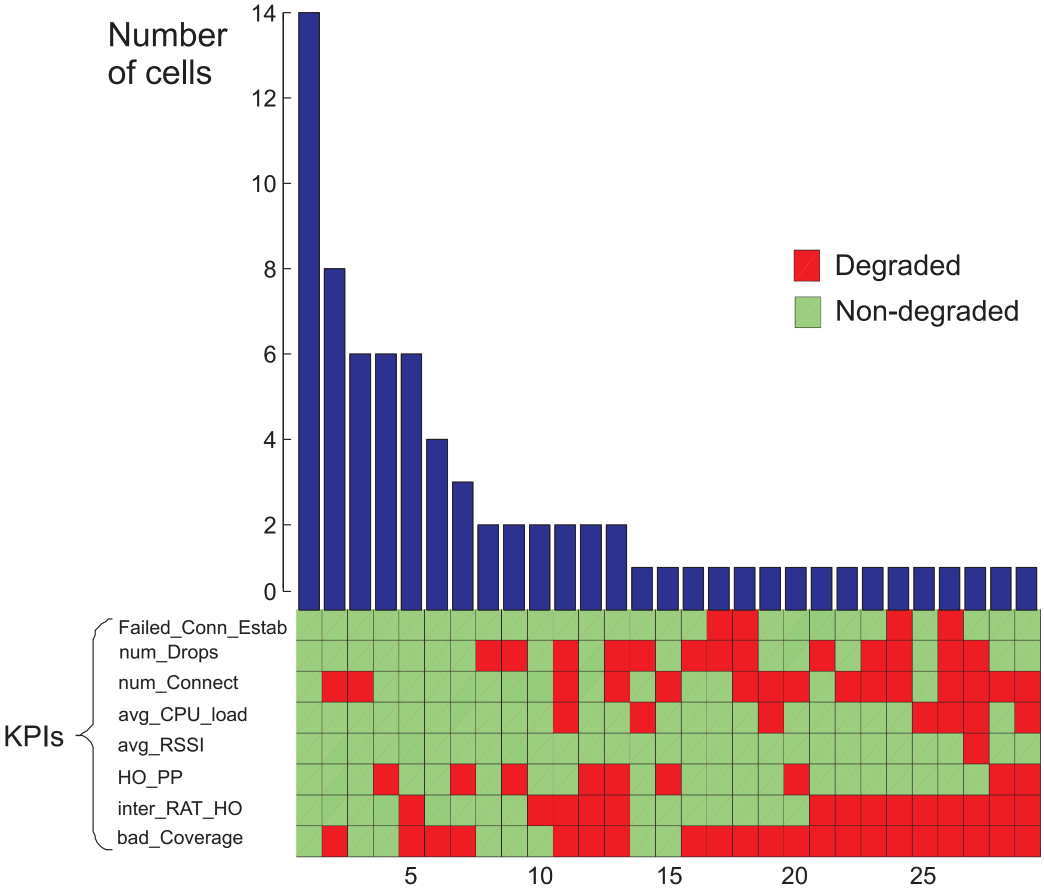

Figure 3: Detected degradation patterns and number of occurrences

27. In these cases, a coverage hole is created affecting both *bad_Coverage* and *Inter_RAT_HO*. In addition, a significant amount of traffic is absorbed by the neighboring cells so that *num_Connect* is increased. The coverage hole and the increased traffic in neighboring cells raises the likelihood of call dropping, which is illustrated by the increase in *num_Drops*. In some cases (see patterns 18, 24 and 26), the cell outage forces neighboring cells to carry an excessive amount of traffic, so that user accessibility is deteriorated. As a result, both *Failed_Conn_Estab* and *num_Connect* will be affected. Most of the remaining patterns are combinations of the previously analyzed cases. Cell outages involving problems of interference in the downlink are more unusual, thus the *avg_RSSI* is rarely affected (see pattern 27).

2.3 A SON approach: an online analysis

It is clear that the higher the number of available data samples under cell outages, the better the analysis is expected to be. However, in a SON context, cell outages occur at the same time they have to be analyzed. This demands for a fast method in order to reduce the traffic losses. In addition, the effectiveness of the compensating actions depends on the success of the analysis. In this section, the proposed algorithm analyzes the cell outage immediately after the detection, so that in general only one sample under the cell outage is available. In particular, two online methods to analyze outage effects have been designed. They are distinguished by the way that degradations in KPIs are determined:

- *Corr_online*: this method calculates the correlation between the observed signal and a reference signal (different from that used in the offline method) which is defined as the unit step function with the jump discontinuity corresponding to the first hour of the cell outage. If the correlation value is above a threshold, the KPI is considered to be degraded. Conceptually, the reference signal represents the cell availability along time. In addition, the length of the signals is shorter than in the offline approach. In particular, only 6 samples are considered (i.e. 5 samples are not affected by the outage and the last sample corresponds to the first hour of the cell outage).

- *Delta_online*: in this method, a threshold is determined as a function of the average and the standard deviation of the KPIs under normal circumpstances. Then, the sample measured under the cell outage is compared to this threshold in order to determine whether the KPI is degraded or not.

The evaluation of these methods has been carried out by using the same dataset than in the offline analysis, which has been taken as the baseline. The results are shown in Table 1, where the second column shows the analysis success rate measured as the number of neighboring cells correctly analyzed divided by the total number of cases. A neighboring cell is successfully analyzed if all their KPIs are correctly identified as degraded or not degraded according to the offline approach. As observed, both methods have similar good performance (above 80% success rate). Thus, these algorithms can be part of effective SON algorithms that quickly reacts to cell outages.

Table 1: Effectiveness of online classification methods

Method	Success rate (%)
Corr_online	81.83
Delta_online	83.17

3 Measuring the load impact on neighboring cells

After a cell outage, the offered traffic in the affected area should be absorbed as much as possible by neighboring cells in the same RAT o another RAT that provides similar user performance. Depending on the particular conditions of each cell outage, the amount of traffic that is lost or steered toward a sub-optimal RAT may be different. In this sense, a neighboring cell with an overlapped coverage area (e.g. by using a secondary carrier frequency) is more likely to absorb a larger amount of traffic than cells with separate coverage areas (e.g. when a frequency reuse of one is used). In the first situation, the compensating actions of the involved cells should be more conservative since they have already partially mitigated the problem. In the second case, the compensating actions should be carried out by cells with partial or limited overlapping coverage areas so that the radio parameter changes need to be larger. Thus, an estimation of the lost traffic can be essential to conduct more rational compensating actions. In this paper, the lost traffic is estimated from KPIs measured in neighboring cells. In particular, the metric *Lost_Traffic_Rate* is calculated as an estimation of the actual lost traffic divided by the total traffic that would be carried by the outage cell in normal conditions. The procedure to calculate this metric is as follows:

- First, the total traffic that would be lost if no traffic was absorbed by neighboring cells is obtained by aggregating the *num_Connect* samples at the outage hours on another day, assuming that the traffic carried during these hours is similar. This term is equivalent to the denominator value of *Lost_Traffic_Rate*.

- Then, the traffic absorbed by neighboring cells is subtracted from the previous term. To do this, the traffic absorbed by each neighboring cell is computed as the difference in *num_Connect* between the outage hours and the same hours on another day, measured in the neighboring cell. Note that the absorbed traffic may lead to call dropping due to bad radio conditions or congestion. For this reason, the *num_Drops* in the neighboring cell should be considered. In particular, the difference in *num_Drops* between the outage hours and the same hours on another day is subtracted from the absorbed traffic to provide a more realistic approximation. In addition, only the main neighboring cells with a positive value of absorbed traffic should be included in the formula.

- Finally, the previous term is divided by the total traffic calculated in the first step. To be consistent, the obtained value is also limited to range between 0 and 1, meaning that the traffic may be totally absorbed or lost by neighboring cells, respectively.

The *Lost_Traffic_Rate* together with the degraded behavior of each neighboring cell obtained in Section 2 has been represented in Fig. 4. Note that the neighboring cells affected by the same cell outage share the same value of *Lost_Traffic_Rate*, since this indicator is related to the cell in outage. As expected, when the impact on KPIs in neighboring cells is none or marginal (i.e. the left part of the figure), the amount of lost traffic is high. Roughly, the higher the number of degraded KPIs, the larger the amount of traffic is expected to be absorbed by neighboring cells. When *num_Connect* and *num_Drops* are affected KPIs, the involvement of neighboring cells in the absorption of traffic is especially important. In addition, when the *avg_CPU_Load* is also degraded, the amount of traffic absorbed by neighboring cells can be even larger, causing important problems in those cells. However, there are some cases for which these assumptions cannot be applied. In particular, some neighboring cells (e.g. patterns 57 and 59) present a non-degraded *num_Connect* and *num_Drops* but the absorbed traffic is higher than expected.

By combining the information provided by the degraded KPIs and traffic lost, the strength of the compensating actions can be better adapted to the specific cell outage.

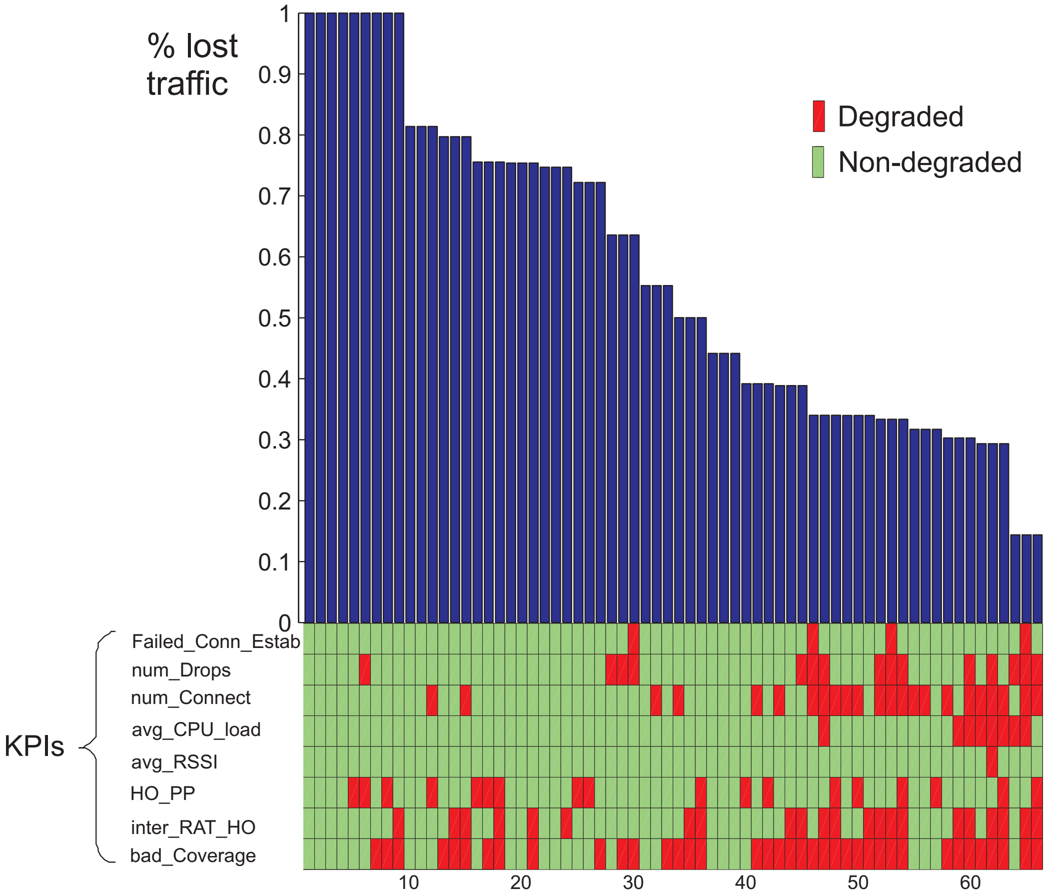

Figure 4: Percentage of lost traffic for each selected neighboring cell

4 Developing advanced COC algorithms

This section discusses some guidelines to improve the existing COC mechanisms by applying cell data analysis. In particular, the proposed structure of COC algorithms should comprise the following phases:

- *Neighboring cell selection*: this stage can be understood as a filter by which neighboring cells are pre-selected to participate in the COC, while the particular role played by each cell will be determined in next stages. The information used to select the most appropriate cells can be related to the network layout (e.g. the degree of overlapping between coverage areas, the location of nodes) or mobility (e.g. based on the number of handover events).

- *Radio parameter selection*: the common effect of cell outages in traditional networks is a coverage hole, which has been typically compensated by tilting the antenna up in neighboring cells. However, if the cell outage involve a lower degree of lost traffic, which has been partially absorbed by the neighboring cells, the tilt should not be adjusted to absorb more traffic. In addition, each cell outage may result in different kinds of degraded KPI patterns that should be compensated in a different way. For example, a high value of *HO_PP* that indicates a waste of radio resources can be mitigated by a proper configuration of the hysteresis parameter. A high value of *inter_RAT_HO* would indicate that many calls are being steered to another RAT (e.g. UMTS) where users may experience poor performance. In this case, the inter-RAT handover margin should also be adjusted to balance the load in a more efficient way. A degradation of *Avg_CPU_Load* may involve congestion. In this case, the traffic load of the neighboring cell should be balanced with respect to other neighboring cells by modifying the intra-RAT handover margin.

- *Determining the magnitude of the compensating action*: in general, SON functionalities progressively adapt radio parameters in order to accomplish a specific goal, which is typically given by reaching certain KPI levels. If the actual values are far from the desired levels, then the actions should be more aggressive, i.e. larger magnitudes of changes are needed. In COC algorithms, the parameter changes in neighboring cells can be modulated according to some data analysis. As explained in Section 3, an estimation of the traffic lost in the affected area becomes a key factor to determine the strength of the COC actions to modify the radio parameters, especially the antenna tilt. For example, if most traffic is inherently absorbed by the neighboring cells, then slighter compensating actions will be required and vice versa.

The last part of the COC algorithm is given by the actual modification of the selected radio parameters until either the stationary state is reached, or the outage cell is recovered. In any case, once the outage cell is recovered, the parameter configuration should be reverted.

5 Conclusions

A novel Cell Outage Compensation methodology has been proposed in this paper. The presented method is based on adapting the compensation strategy according to the degradation produced by the cell outage in the neighboring cells. The adaptation consists of automatically selecting the compensating cells, the configuration parameters to be modified and the magnitude of the compensating action. A study of the degradation produced by cell outages in the neighboring cells in a real network has been presented in this paper. Specifically, three methods have been proposed. One of them aims to analyze the degradation produced by the cell outage in the neighboring cells based on KPIs correlation using historical records of cell outages. Conversely, in the context of SON, the two other proposed approaches are executed immediately after the outage detection. The first one is based on correlations and the other is based on delta detection. Both methods allow to determine the type of cell outage and, thus, to effectively adapt the compensation algorithm. In addition, an estimation of the lost traffic caused by the cell outages has been obtained. The results have shown that cell outages may result in different issues which are reflected by different types of degradations in the neighboring cells. Such a set of different situations should be compensated in a different way.

Acknowledgment

This work has been partially funded by Optimi-Ericsson, Junta de Andalucía (Agencia IDEA, Consejería de Ciencia, Innovación y Empresa, ref. 59288; and Proyecto de Investigación de Excelencia P12-TIC-2905) and ERDF.

References

[1] "3GPP, Telecommunication management; Self-Organizing Networks (SON); Concepts and requirements," 3rd Generation Partnership Project (3GPP), TS 32.500, 2014.

[2] A. Imran and A. Zoha, "Challenges in 5G: how to empower SON with big data for enabling 5G," *Network, IEEE*, vol. 28, no. 6, pp. 27–33, Nov 2014.

[3] R. Barco, P. Lázaro, and P. Muñoz, "A unified framework for self-healing in wireless networks," in *IEEE Communications Magazine*, Dec 2012, pp. vol.50 (12), pp.134–142.

[4] "3GPP, Telecommunication management; Self-Organizing Networks (SON); Self-healing concepts and requirements," 3rd Generation Partnership Project (3GPP), TS 32.541, 2014.

[5] I. de la Bandera, R. Barco, P. Muñoz, and I. Serrano, "Cell outage detection based on handover statistics," *Communications Letters, IEEE*, vol. 19, no. 7, pp. 1189–1192, 2015.

[6] O. Onireti, A. Zoha, J. Moysen, A. Imran, L. Giupponi, M. Imran, and A. Abu Dayya, "A cell outage management framework for dense heterogeneous networks," *Vehicular Technology, IEEE Transactions on*, vol. PP, no. 99, pp. 1–1, 2015.

[7] L. Kayili and E. Sousa, "Cell outage compensation for irregular cellular networks," in *Wireless Communications and Networking Conference (WCNC), 2014 IEEE*, April 2014, pp. 1850–1855.

[8] J. Moysen and L. Giupponi, "A Reinforcement Learning based solution for Self-Healing in LTE networks," in *Vehicular Technology Conference (VTC Fall), 2014 IEEE 80th*, Sept 2014, pp. 1–6.

[9] M. Asghar, R. Fehlmann, and T. Ristaniemi, "Correlation-based cell degradation detection for operational fault detection in cellular wireless base-stations," in *Mobile Networks and Management*, ser. Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering, D. Pesch, A. Timm-Giel, R. Calvo, B.-L. Wenning, and K. Pentikousis, Eds. Springer International Publishing, 2013, vol. 125, pp. 83–93.

[10] B. Cheung, G. Kumar, and S. A. Rao, "Statistical algorithms in fault detection and prediction: Toward a healthier network," *Bell Labs Technical Journal*, vol. 9, no. 4, pp. 171–185, 2005.