November 2016

SMALL CELL REAL-TIME ANALYTICS

CASE STUDY

Table of Contents

Glossary

QoE	Quality of Experience. The quality of connection as perceived by the end user of a Communications Service Provider.
QoS	Quality of Service. The network ability to achieve maximum bandwidth and deal with other network performance elements sucn as latency, error rate, and uptime.
WiFi	A technology by which electronic devices connect to a wireless LAN through an unlicensed radio spectrum.
CSP	Communications Service Provider.
SON	3GPP. Self-Organizing Networks. An automation intelligent technology designed to make the planning, configuration, management, optimization, and healing of mobile radio access networks simpler and faster using machine learning algorithms.
3GPP	The 3rd Generation Partnership Project unites telecommunication standard development organizations and produces the Reports and Specifications that define 3GPP technologies.
NFV	Network Function Virtualization. A network architecture concept that uses the technologies of IT virtualization to virtualize entire classes of network node functions into building blocks that can connect, or chain together, to create communication services.
RAN	Radio Access Network. A set of radio technology, which connects mobile devices with the core network using licensed radio spectrums.
C-RAN	Cloud RAN or Centralized RAN. Provides centralized baseband processing for many cells is. Benefits of C-RAN include improved performance due to the ability to coordinate between cells, and cost reductions as a result of pooling resources.
HetNets	Heterogeneous Network. It is typically composed of multiple radio access technologies, architectures, transmission solutions, and base stations of varying transmission power.
OSS	Operations Support Systems. The support management functions such as network inventory, service provisioning, network configuration and fault management.
NOC	Network Operations Centre.
KPI	Key Performance Indicator.
MDT	Minimization of Drive Tests
CDR	Charging Detail Record

Executive Summary

Due to the increase in smartphone functionality and the growing number of interconnected heterogeneous devices, telecommunications networks are experiencing an unprecedented explosion in the volume of data.

In today's mobile networks, approximately 80% of mobile data traffic is generated indoors. Small Cells are supposed to be plasters, little boosts of low-power spectrum, which close gaps in the radio coverage of macro cells and solve densification in a cost-effective way. However, despite the initial investment for Small Cells being relatively low, the numbers in which they are needed requires a significant investment, which needs serious operational strategies.

Operators cannot afford to roll trucks for each small cell deployment. They need to master the techniques for applying real-time analytics and incrementing the self-operating network capabilities, paving the road for the adoption of Self-Organizing Networks (SON).

The importance of advanced analytics in the successful adoption of SON is unanimously accepted by the industry. Because SON seems to be the only feasible solution for such formidable network operational challenges, analytics are a top investment priority for some telecom operators around the world. However, despite the substantial investment in creating internal data science groups, many organizations struggle to establish a sustainable data science capability able to embed data science in their network products and generate the promised return on investment. There are several reasons for this:

> Data Scientists are not network architects. The knowledge transfer between domain expertise and data science is not always easy. Statistical models bring mathematical truths, but the correct interpretation and its application to the business is not trivial and requires knowledge in both areas.
>
> Data Scientists have a high rate of professional mobility. Data Scientists are very well-recognized professionals and are paid accordingly. Retaining talented Data Scientists is expensive and requires additional investment with an uncertain return, which increases the scepticism from leadership.
>
> Organizational issues. Data Science groups are cross-functional teams who usually lack specific domain knowledge. This often generates resistance in other organizations who do not recognize the value of data science and usually consider data scientists as outsiders.

Thingbook.IO is a streaming predictive analytics company designed to process massive amounts of data. By using the Thingbook.IO machine learning capabilities, CSPs can predict service degradation, automatically discover Cell Outage, Sleeping cell, QoE degradation, and Cell anomalies in real time.

This case study illustrates the results obtained by applying Thingbook.IO Turing (Stream-based prediction/analysis engine) to the problem of operating massive Small Cells in a dense area of a large European city. Thingbook.IO Turing was proven to be successful in detecting cell outages, KPI anomalies, and isolating the subscribers impacted by poor QoE. This study analyzes the rate of churn and correlates it with the subscribers systematically affected by cell outages and failed hand-over process.

Network Challenges and the Role of Small Cells

Network Situation

Originally intended for making calls, the cell phone has become the most widely used piece of high technology, supporting a variety of multimedia services. As such, smartphones are an integral part of our everyday life. The intensive usage has not only raised the volume of network traffic, but has put strict Quality of Experience (QoE) demands to meet user expectations.

Users use smartphones to write emails, watch videos, and even make video calls. With the growing number of mobile phone subscribers, the number of media services increases. At the same time, end-users are becoming more and more demanding regarding the quality of service provided. According to Ericsson, these factors will lead to the exponential growth of mobile traffic consumption over the coming years.

Facing this new network reality, any consideration of a mobile network build must begin with data volumes. Data drives all network decisions, from spectrum allocation to backhaul bandwidth, and from site selection to technology type. The growing volume of data and the intelligence derived from it, has a decisive influence on any aspect of the network. The densification of a network structure with small cells is regarded as a key solution to meet growing capacity demands, while keeping the level of investment low.

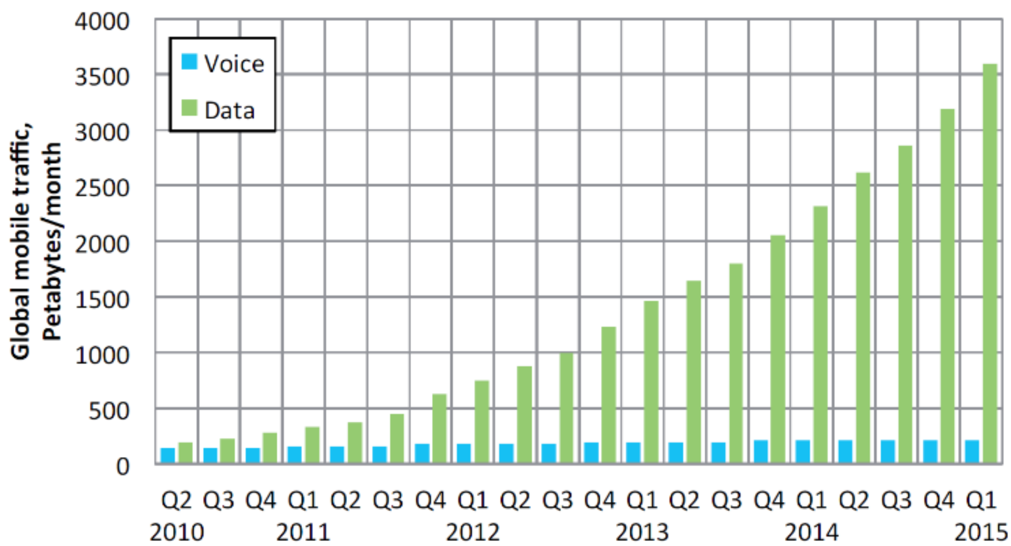

Figure 1. Shows the growth in data usage over the last five years.

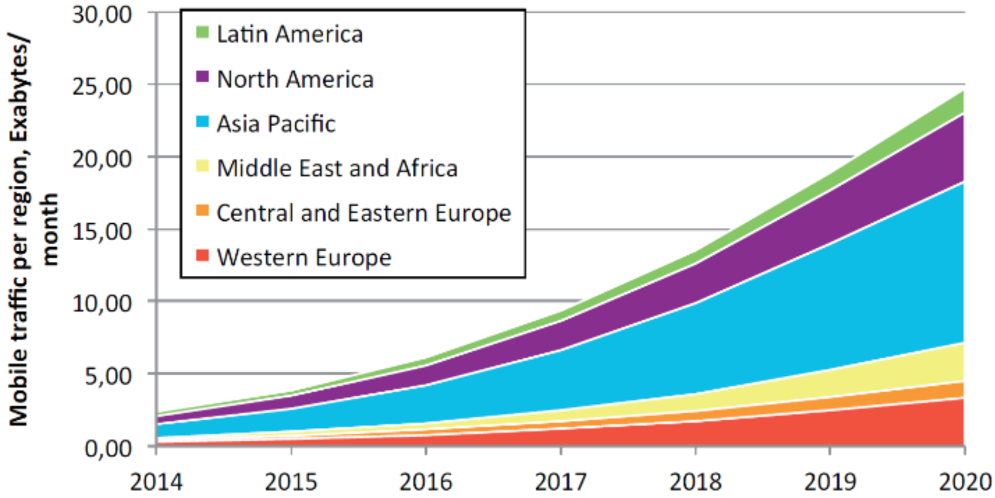

Figure 2. Shows the latest projection for data volume growth over the next four years. Such growth in network traffic is unprecedented.

Small Cells and HetNet Come to the Rescue

In today's networks, approximately 80% of mobile data traffic is generated indoors. Small cells were intended to be plasters, little boosts of low-power spectrum, which close gaps in the radio coverage of macro cells and solve the densification problem.

The macro cells were expected to continue their role as the pillars of the mobile network, and require all the support and exactitude of design worthy of their station and scale. However, as small cells proliferate and combine to create heterogeneous networks (HetNet), it becomes more evident that their role in delivering mobile connectivity extends far beyond the initial expectations. For mobile operators today, Small Cells are the only cost-effective way to keep pace with market growth and customer demand.

As Cloud and Network Functions Virtualization (NFV) begin to play a role, software tools for planning, automation and assurance that specifically target small-cell deployment, will also get additional attention and demand. The challenge for mobile operators is not the decision of which solution to select, but how to best integrate multiple technologies within their networks, how to find the right balance to maximize their cumulative benefits, and how to leverage existing assets to facilitate the evolution of their networks. Operators need a roadmap for a long-term and multi-step strategy to increase RAN capacity density in high-traffic areas to meet the subscriber demand for data services.

Primary Operational Issues with Small Cell Implementation

Despite the initial investment for Small Cells being relatively low, the numbers in which they are needed requires a significant investment, which needs serious operational strategies. The traffic that Small Cells can deliver and offload from the larger radios is an important factor in meeting customer growth and the expected Quality of Experience. Nevertheless, the collected data concerning customer usage and service assurance are equally important, beneficial, and strategically important.

CSPs cannot afford to roll trucks for each small cell deployment. This operational reality has driven the need for automation in the turn-up process, and for new modeling processes that need to be completed long before the automation of SON takes over.

These operational challenges are not much different to those for macro cells. However, scaling up on the number of radios and scaling down on the cost structure, only serves to exacerbate many of these challenges. For instance, deploying small cells using traditional truck rolls is a non-starter from a cost perspective. Such is also the case with ongoing maintenance, as technicians cannot cost-effectively make manual adjustments when interference or directional issues arise, especially because they will be in far greater numbers.

Another inherent operational and integration issue lies with HetNets and the handoff to WiFi networks. The technologies have progressed down separate paths that go beyond their use of licensed versus unlicensed spectrum. CSPs have little, if any, visibility into WiFi hotspots or which hotspot a device might be communicating with at any point in time. WiFi devices "see" many access points, and often switch connectivity back and forth. Because of this lack of visibility, there has not been much in the way of auto-discovery capabilities for WiFi hotspots.

Automation and Self-Organizing Networks Seem to be the Answer

Mainly motivated by the complexity introduced by Small Cells and HetNets, 3GPP introduced Self-Organizing Networks (SON). The SON concept proposes a set of situations or use cases in which the network can be self-configurated, self-optimized, and self-healed. However, SON and automation are not the same thing. SON Products have proven difficult to adopt by the CSP, because they are not yet ready to entrust their networks to autonomy.

Automating processes and parameter changes in a network comprised of a single vendor's RAN technology is difficult enough, but HetNets are the opposite scenario to a single-vendor network environment. Not only HetNets will be affected by the radio gear from multiple vendors, various forms of backhaul, and present the mentioned integration challenges, but the required machine learning algorithms that determine the next best action in all performance and customer scenarios are still to be written.

So far, CSPs have not mastered the techniques for applying real-time analytics, particularly the correlation of network and non-network data to identify and predict service degradation situations. Network and non-network data is required to make real-time autonomous decisions that serve the overall business, and not just some estimated performance benchmark.

In the multi-technology, multi-vendor, multi-frequency environment of HetNets, which are in turn managed and monetized by diverse ecosystems of support software providers, integration is complex. In fact, integration issues still pose a threat to the success of originally envisioned HetNet and C-RAN solutions.

Many aspects in a RAN can be automated, and already are automated to a limited extent. Antenna tilt has been automated in some cases, as has the adjustment of power levels and receiver gain. The ability to download configurations has also received a level of automation; radios can now put themselves to sleep to save power. However, most configuration or parameter changes are not done globally and, when it comes to small cells, re-configuration must be done with the macro network in mind.

The main role of Small Cells is to serve the macro cell. A very small number of use cases have been designed and tested to the degree necessary to build analytics and automation with the entire picture in mind. There are many shared resources between the two, as well as potentially between competing CSPs that share radio resources with others. This heterogeneous environment requires reliable, real-time data analytics systems for dynamic configuration management, service assurance, and network monitoring.

QoS and assurance are also difficult to manage because of the variety of traffic types traversing the network, and the lack of performance monitoring visibility into some Small Cells. The following are examples of where automation can be applied to the Small Cell implementation process without the need to broach the subject of full SON:

- CSPs can incrementally experiment with more advanced machine learning techniques, which can predict likely service degradation, as well as the impact of applying correction actions generated dynamically.
- CSPs can also apply machine learning analysis to include dynamic optimization, self-healing, and closed loop operations management.

It should be possible to localize any potential issues caused by automation in a small cell, and prevent it from cascading throughout the network.

Thingbook.IO Telecom Analytics. Automatization and SON

Introduction to Thingbook.IO Network Analytics

With the buzz around big data technologies and infrastructures, many CSPs struggle to extract actionable insights from their network data. In this sense, the main purpose of ThingBook.IO is to allow data-driven decision-making processes in contexts dominated by strong volume and velocity requirements like telecom networks.

ThingBook.IO applies stream machine learning to extract sufficient information to detect anomalies, correlate behaviors and Key Performance Indicators (KPIs), discover new behavioral patterns, and perform root-cause analysis qualifying the CSP to make informed decisions in real-time.

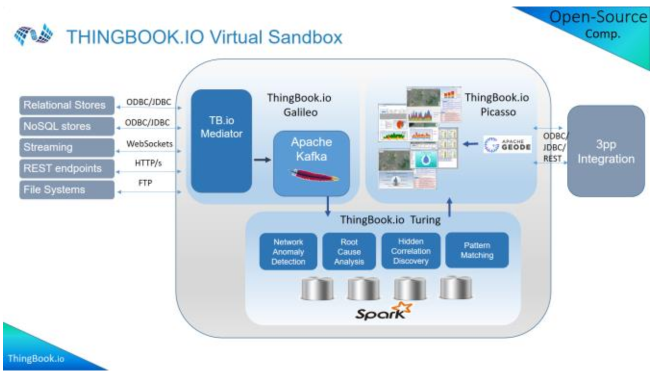

Thingbook.IO Sandbox Components

This real-time capability allows both humans and expert systems to take preventive or corrective measures.

All ThingBook.IO actionable analytics applications are built on proven experience of ThingBook.IO in Cellular Networks, Machine Learning, Streaming technology, and actual use-cases from the telecommunication industry, to answer questions such as:

- How do I analyze customer experience and reduce care costs by anticipating customer needs?
- How do I optimize my network investment to focus on the experience of high valued customers and reduce significant churn?
- How do I improve operational processes and reduce operational expenses?

To do so, ThingBook.IO analyzes large volumes of data, which are in turn processed in real-time to: (1) perform data harvesting and mapping (2) perform data profiling and in-depth analysis and (3) provide visualizations for human consumption and integration with external systems.

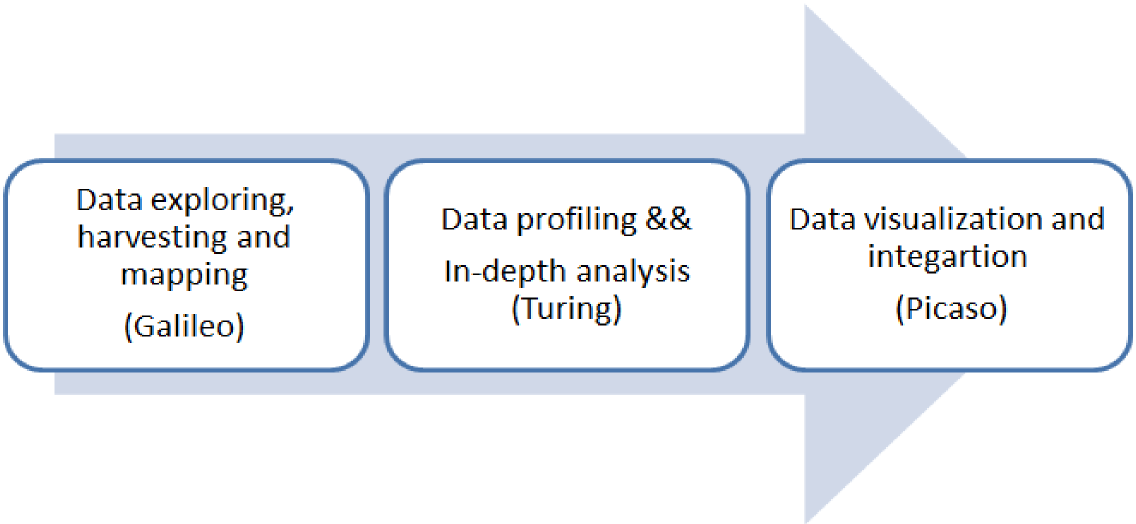

Data harvesting (Galileo)

Data retrieval and initial processing is usually a time-consuming task, prone to errors and human intervention. ThingBook.IO implements a specific subsystem (also known as Galileo) for efficiently collecting, mapping, aggregating, and moving large amounts of data from many different sources to a centralized data processing point.

Galileo is not restricted to a single type of data, such as packet traces, Network databases, and files. Different data sources can be simultaneously used and customizable according to the specific requirements of the deployment environment.

In-Depth Analysis (Turing)

Turing uses machine learning techniques to conduct real-time analysis of the data. It implements the core capabilities of ThingBook.IO Network products: Anomaly detection, Pattern Matching, Probable Cause Analysis and KPI dependencies, and performance ranking.

Network Anomaly Detection

Network Anomaly Detection is the action of finding behaviors in network traffic, often short-lived, which do not conform to expected patterns. These non-conforming behaviors find extensive use in a wide variety of telecommunication networks, such as QoE impact, equipment performance degradation, security and intrusion detection or attack blockers when the anomalies are detected in the early stages.

Turing includes the ability to proactively detect network anomalies and detect unknown network behaviors without using any signatures, labeled traffic, or training. It uses an unsupervised outlier detection approach, based on Sub-Space Clustering and Multiple Evidence Accumulation techniques to pinpoint different kinds of network anomalies.

Behavioral Matching

A pattern can be considered as the mathematical expression of specific behavior. Such behavior can either correspond to a newly discovered knowledge or something learned in the past. In either case, Turing has the capability to store and recognize meaningful data behaviors in real-time.

The ability to recognize behaviors in the data has a tremendous implication in the way Operations Support Systems (OSS) detect pre-defined network incidents, such as cell congestions, cell outage, or sleeping cells.

In current network systems, by monitoring certain thresholds in various Network Elements (NEs), warnings and alarms are triggered to domain experts in the Network Operations Centre (NOC), which indicate an incident in the network. The domain experts then manually investigate the reports starting with the most critical ones. The process of resolving nontrivial network incidents requires significant comprehensive knowledge about the network architecture, its elements, and their capabilities.

Turing uses a patented unsupervised approach to measuring distances between any individual entity (such as a cell, or a KPI) and behaviors either previously identified and labeled as abnormal, or automatically learnt as a deviation from the normal expected behavior. Turing detects the gradual approach of a deviation from the expected behavior and alerts the NOC about the new situation. This method avoids long post-mortem investigation times.

Root Cause Isolation

Root Cause Isolation (RCI) is the process of identifying the source of anomalies (potentially problems) in a system using only data observation. Many OSS systems and NOCs suffer from a common problem: when the network fails to function correctly, it is often difficult to determine which part is the source of the problem.

The fundamental challenge is that the symptoms of a failure often manifest as end-to-end failures in the operation of the system, without causing obvious failures in the system components; noticing that something has gone wrong does not necessarily provide information about where to look to fix it.

Turing learns one or multiple abstract models for the normal network operation by capturing the commonalities and dependencies between KPIs, traffic structures, and network equipment. By comparing the learned models with the failure situation in real-time, Turing provides intelligence of the components leading the failure.

KPI Dependencies and Performance Ranking

Deploying dedicated quality assurance systems has become vital for CSPs to guarantee that the connectivity services are provided with the highest possible quality. The complexity of such monitoring systems grows with the complexity of communication services and with the demand for providing complex monitoring insight in real-time.

In such complex systems, the nontrivial relation between different KPIs, as well as an erratic behavior of any KPI across different cells, areas, or equipment, can mask a configuration problem or highlight a poorly covered area.

Turing addresses the problem of analyzing a KPIs behavior by discovering its dependencies and determining non-justified changes on the metrics values. Turing uses KPI behavioral models to determine cell-performance status and can cope with legitimate system changes, that is, concept drift.

Data Visualization and Results Integration (Picasso)

Information consumption processes arguably benefit from user friendly interfaces, preferably with graphical capabilities and rich visualizations. In this sense, because Picasso operates in a dimensionally reduced space, it is possible to keep complexity to a minimum.

In today´s interconnected world, there is no room for stand-alone systems. Any single functionality represents a fraction of a higher and more complex function that requires integration and orchestration. Picasso uses different communications and messaging, such as Kafka, SNMP, and database services, to integrate the generated results into external systems.

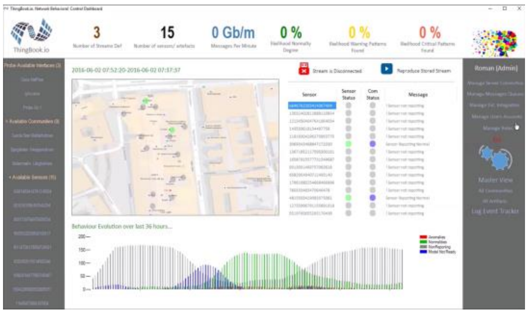

Thingbook.IO and Small Cells Analytics

In today's data-saturated telecommunications networks, correlating customer perceptions of quality with the network performance data has become a formidable challenge for CSPs around the globe.

This document describes how Thingbook.IO Machine Learning capabilities were applied to proactively find the entities (areas, bands, Cells, Times, and subscribers) performing differently from the normal or expected behavior.

The main contribution of this work is the application of a streaming behavioral analytics framework (Thingbook.io Turing) using data from Small Cells KPIs, MDT (Minimization of Drive Test) data and Billing system data to implement an analytics framework that automatically detects impaired cells, anomalous KPI behaviors and the root-cause of the anomalies. Such automatic functionality enables mobile operators to increase network robustness and decrease maintenance costs.

Cell Outage Detection Mechanisim

Proactive detection of malfunctioning cells is part of the self-healing mechanism. In general, cell outage takes place due to multiple reason, such as hardware or software failures (including missconfigurations or sotware bugs) or even environmental changes. Usually, the detection of a malfunctioning cell is performed through the analysis of alarms, KPIs, or in many cases, multiple customer complaints. Generally, cell outage is classified into three types: degraded, crippled, or catatonic (Cheung, Fishkin, Kumar and Rao in their work "Method of monitoring wireless network performance”).

Due to its complexity and ambiguous behaviour, the most difficult to detect is a degraded cell. A degraded cell can carry network traffic, but not as much as a correctly functioning cell. A crippled cell is characterized by severely degraded performance, but still provides a service to a few users. A cell which experiences complete inoperability is referred to as a catatonic cell.

A sleeping cell is a cell degraded type, which is invisible for network operators through traditional alarms. This peculiarity makes a sleeping cell problem a very challenging task. Usually, this type of a cell outage becomes visible after drive tests, using the detailed analysis of KPIs, or because of end user complaints. However, the cellular network still services the mobile terminals, which were synchronized with the broken cell before the problem occurred. Thus, with full signal coverage there is a lack of connectivity. This characteristic of Random Access Channel failure makes it a challenging problem.

According to this specification, random access procedure is deployed for initial access, radio resource connection protocol re-establishment, uplink synchronization, and handover. When RACH failure occurs, it prevents a user device from establishing a connection or making a handover to the malfunctioning cell. However, the problematic cell still carries traffic for those users who were connected before the failure occurred. Because, only a few users notice the problem the cell belongs to an impaired outage type. Due to mobility, users leave the malfunctioning area and the cell becomes catatonic.

Usually, RACH problems are detected only after a long observation time or after multiple user complaints. Thus, a timely detection of RACH failures is an important issue in LTE networks. The longer the detection takes, the higer the possibilities of churn. In today's networks, the detection of a sleeping cell relies on the manual analysis of KPIs, because traditional alarms are not raised. It is a tedious, error-prone and expensive process. (RACH is a constitutional part of LTE uplink transport and physical levels, see (3GPP, 2015a)).

This study illustrates the framework accuracy in detecting malfunctioning cells, which are invisible for network operators using traditional alarms, the subscribers affected by malfunctioning and its correlation with the churn levels. It is proven that subscribers affected by performance degradation problems have a significantly higher level of churn than those who were not. Thingbook.IO Turing capabilities (explained previously) addressed the problem of automatic cell outage detection and root cause analysis in LTE networks.

Turing implements an anomaly detection and diagnosis framework for LTE networks. The underlying procedure is the analysis of channel quality degradation through cell-wise profile behaviors of network KPIs and MDT data. If the deviation of a KPI profile from normal behavior is detected, then an alarm is raised. The deviation might be caused by the degradation of cell performance or a cell outage. Therefore, even small deteriorations can be identified. When the framework detects a malfunctioning cell, then root cause analysis is performed and the subscribers in the area covered by that cell and neighboring are tracked.

Root-Cause analysis involves an automatic investigation of problem KPIs and diagnosis regarding failure reasons. Turing fully automates the diagnosis process by creating models per cell, KPI and area to identify the component leading the anomaly. To overcome the problem of spurious data, Thingbook.IO implements a cooperative scheme for cell outage detection where neighboring cells share information. Each cell categorizes the coming flow of information from surrounding units. The behaviour characterization of the community of cells using the hand-over events and subscriber movement patterns, provides higher accuracy than traditional methods.

Minimization of Drive Tests Data

Drive tests are field measurements performed by mobile operators and aimed at gathering network statistics about coverage, capacity, and quality of service. Drive tests discover the weak parts of the network and provide the means for performance improvement. Therefore, field measurements are essential for network planning, optimization, and troubleshooting processes. Drive tests are carried out by means of special vehicles transporting the measuring equipment along the streets and roads. As such, mobile operators obtain coverage maps and network conditions as a function of coordinates.

Despite its benefits, a drive testing campaign has noticeable drawbacks. Drive tests do not provide full radio signal measurement coverage: pedestrian zones, buildings, and other car inaccessible spots are out of scope. As a response, 3GPP introduced MDT functionality in Release 10 (3GPP, 2014a), addressing the high demands of network operators for the automation of driven tests. The underlying idea is to employ user handsets for the collection of network performance information. Unlike traditional field measurements, MDT can collect network data throughout the map, including indoor and outdoor locations.

MDT architecture is built on a subscriber and equipment trace concept. Tracing can keep track of all user activities in a cellular network. 3GPP also distinguishes two ways of tracing activation: signaling based and management based. Signaling based trace sessions monitor all activities of a chosen handset. Monitoring continues if the handset moves throughout the whole network area. Management based tracing captures all users within a chosen area of several cells. This type of tracing is activated when a new user comes into the specified area and is interrupted when a user leaves the area.

Data Specification Used in the Study

Thingbook.IO uses the follow data sources:

- For each cell, 12 KPIs collected every 60'' for four months, from 15/9/2014 to 15/01/2015. The KPIs have different characteristics: some of them are measurements of user traffic utilization (such as downlink or uplink data volume or throughput), while others are measurements of call control parameters (such as drop-call rate and successful call-setup rate).
- MDT Management data captured in intervals of 30'' in the dense area covered by the study for 3 months. Used Data Specification can be found at http://www.etsi.org/deliver/etsi_ts/132400_132499/132423/08.01.00_60/ts_132423v080100p.pdf
- 6 months of CDR data. CDR data structure can be found at: http://www.3gpp.org/ftp/tsg_sa/WG5_TM/TSGS5_68/_specs_for_checking/32298-920.doc

Study Highlights

Strong behavioral differences across the population of cells. Top 12% of the cells represents 60% of the total activity. Last 8% of the cells represents 0.010% of the activity.

Studied Cells show a strong seasonality factor describing substantial behavioral differences depending on factors such as day and time. Seasonality does not affect all the cells equally. Due to the seasonality factor and the disparity of its impact, the global behavioral model provides suboptimal results.

An important factor for data-driven approaches is the amount of available data. During night-time the number of users is significantly smaller than during the daytime in a city. It represents a lack factor for the detection of antenna gain failure.

Result Highlights

Six groups of cells were identified based on behavior similarities (Top 5 concentrated the 51.6% of the cells)

1.3% of the studied behavior was found abnormal or strongly different to the expected learned model.

Two Cells were repeatedly found in “sleeping/outage” status affecting a total of 28739 subscribers over the time the proof of concept was conducted, experimented connections refused and poor QoE.

The handset model and OS implemented have a strong influence on the call quality. The number of CS Call Setup Success Rate and CS RAB Establishment Success Rate decrease systematically for the most used model of handset vs. the second most used model.