Application of AI/ML Technology to Address Congestion, Quality, and Security of Private Military Network Deployments

Dr. M. Patryk Debicki
Thales UK Limited
UNITED KINGDOM

patryk.debicki@thalesgroup.com

Gonzalo Aréchaga
Thales Programas de Electrónica y Comunicaciones S.A.U.
Diego Marín Aguilera 1, 28919 Leganés, Madrid
SPAIN

gonzalo.arechaga@thalesgroup.com

ABSTRACT

With the quick evolution of mobile networks in public domain, military Communications Systems increasingly adopt and enhance what's already available in the commercial space. There is a growing number of tactical LTE deployments that provide improved bandwidth and latency when compared to legacy non-3GPP based systems. With private 5G deployments around the corner, more network improvements are coming along that will enable latency-critical applications to be controlled over the air. However, all this improvement cannot be achieved without increasing complexities of networks. Consequently, traditional network management and operation systems need to be augmented with Machine Learning (ML) and Artificial Intelligence (AI) technology as it becomes physically impossible for humans to analyse the ever-increasing amount of information coming from different layers and domains on a modern cellular network.

Current deployment of communications on military operations is usually based on military-only communications. In some cases, commercial networks with ciphering devices to secure them are also used. The wide deployment of 4G networks and the future deployment in the short term of 5G technology would be a game changer. This high-band mobile networks would enable both military-only networks based on these technologies and hybrid deployments where bubbles of military devices are connected through commercial networks. As an example, a platoon using mobile terminals would exchange information between them on a private network, communicating with a second private network of Internet of Things sensor devices on the battlefield, and sharing the needed information with a third private network on the Command Post, while these three networks are interconnected via a commercial network, all of them using the same technology.

In this paper we are going to talk about AI/ML based solutions that improve the efficiency of network operations mixing private and public networks by detecting anomalies in real time and predicting potential future issues. We will present use cases that address congestion, quality, and security related concerns (information leaks or a security attack) and show to efficiently identify and prevent upcoming problems. We will also talk about the evolution into 5G systems and how the AI/ML technology is inherently embedded within the 3GPP standards and how it can help in tactical military 5G network deployments.

1.INTRODUCTION

Over the years the mobile core network kept evolving to embrace new capabilities being developed in radio network, devices, cloud technology as well as by the requirements of new services. With the vast advancement in mobile communication, it started to make sense to adopt and enhance what is already available in the commercial space. In fact, with the significant improvements in throughput and latency made in LTE (Long Term Evolution, aka 4G) technology (compared to its 2G and 3G predecessors), Thales has already been delivering different flavours of private LTE bubbles that would utilise commercial transmission enhanced with Thales security technology sitting on top. With first 5G Stand Alone (5G SA) networks coming up and with embedded end-to-end slicing capability and standardised analytics functions, new opportunities are emerging when it comes to the flexibility of a bubble deployment, support for new use cases and the ease of the network operations.

From a military perspective, the mix of private and commercial networks provides multiple advantages, but also some disadvantages. A private network is per definition under the control of the owner, which can provide more security for its usage, but the cost of deployment is usually significantly higher. On the other hand, the use of public networks provides a lower cost with a lower security and greater risks. The mix of both worlds can optimize the use of resources in military applications. Bubbles of pure private networks can be interconnected via commercial networks with additional security layers, and even more, 5G SA (Stand Alone) via network slicing can even allow a low-cost deployment.

In order to cope with the associated security issues that could arise in these scenarios, we propose the use of the following network management techniques thanks to the use of Machine Learning (ML).

2 NETWORK OPERATIONS

Traditional Performance Management

Traditional performance management systems were based on tens of thousands of metrics, thousands of Key Performance Indicators (KPIs) defined for each network domain (you could say that is an oxymoron as KPIs should only be few), hundreds of reports, tens of employees engaged in creating and analysing those reports based on multiple disparate reporting systems and all of it was enwrapped by a manual process prone to missing important information. This approached is presented in Figure 1.

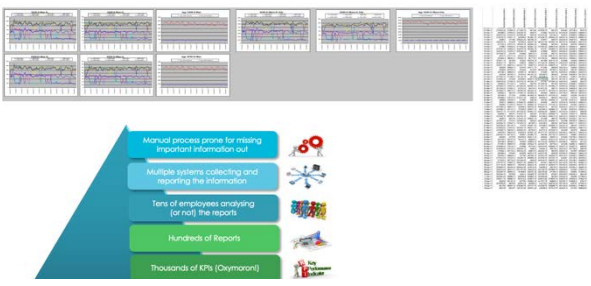

Figure 1. Traditional approach to performance management.

While this approach worked fairly well in the past with well established processes in the Communication Service Providers (CSP) organisations, with every new network generation and new services and capabilities being added in, this approached proved to be increasingly complicated and unfit to manage 3rd generation network. With LTE deployments many CSP started working on different approach and started working on

so-called NOC (Network Operation Centres) to SOC (Service Operation Centre) transformation in order to get a better visibility of service quality delivered to the end customers. It required heavy investments and long process and, in many cases, these transformations were not as successful as the CSP had expected. SOC systems indeed provide much better visibility of user experience, however, the troubleshooting process is still quite complicated and rarely is integrated with alarm-management systems or trouble-ticketing systems.

Performance Management Augmented with ML

An alternative approach used by Thales provides a powerful capability to deliver a complete picture of the overall network performance, considering all network domains and raw event-based data.
The approach is based on advanced anomaly detection algorithms and ensemble functions that allow for cross correlation of the anomalies over different domains. Consequently, it allows for automated discovery of important issues and anomalies on the network and the triggering of an action in the form of a performance-based alarm notification to a customer-care system or perhaps even changing the network policy. This approach is depicted in Figure 2.

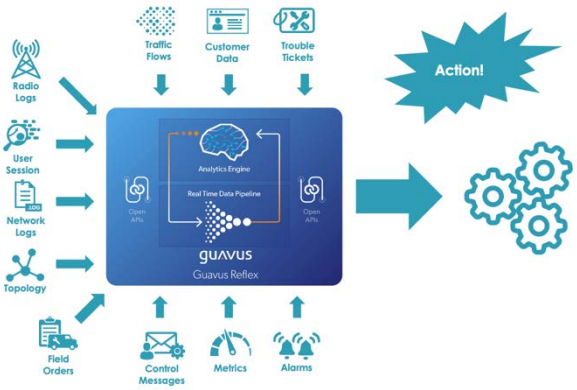

Figure 2. Alternative approach to performance management that is based on anomaly detection and cross-correlation.

Anomaly Detection Algorithms

In standard deployments anomaly detection is threshold based, where such thresholds exceed a computed moving standard deviation metric of past distributions of a value or distributions of values in specific time ranges (such as peak and off-peak times or weekday or a weekend). However, due to the variety of measures and types of metrics that are monitored in network operations, such an approach is not always sufficient to identify a number of different types of important anomalies. We have implemented several methods to capture such anomalies that are not otherwise detectable by the usual standard monitoring approach.

Figure 3 provides a view of types of anomalies that typically are encountered in a telecom network.

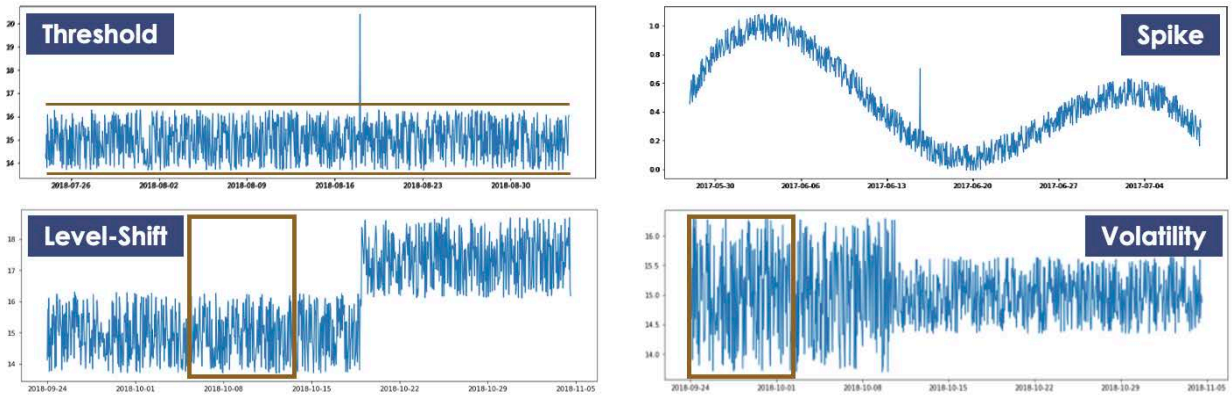

Figure 3. Anomaly Detection Types.

All of the algorithms to detect these anomalies are based on a standard score (or Z-score – its absolute value represents the distance between that raw score and the population mean in units of the standard deviation; the value is negative when the raw score is below the mean, positive when above) where calculations are applied in a different way:

- **Threshold** computes Z-scores for data points in the raw time series
- **Spike** generates a time series for difference of consecutive raw data points, then computes Z-scores
- **Level-Shift** generates a time series for mean of raw data points over a sliding window, then computes Z-scores
- **Volatility** generates a time series for std dev of raw data points over a sliding window, then computes Z-scores

Another example of an anomaly detection algorithm is presented in Figure 4. It is based on prediction algorithms that allow us to identify anomalies in seasonal patterns.

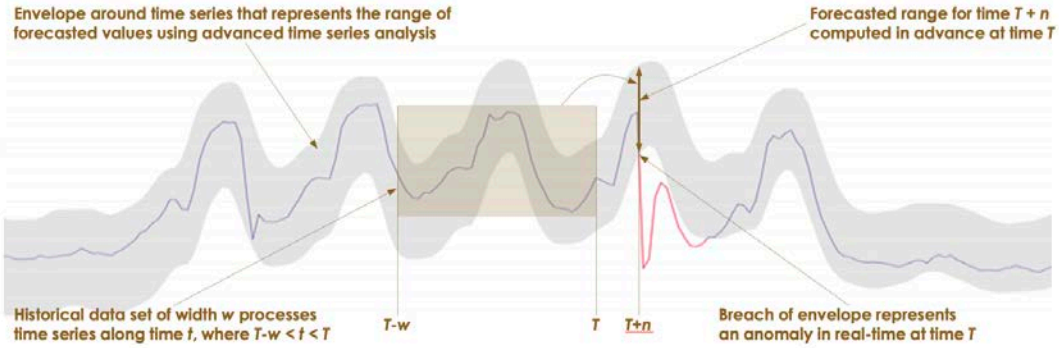

Figure 4. Season anomaly detection algorithm.

Correlating Multiple Domains

Identification of an anomaly in different types of counters on the network, does not necessarily identify a method of resolving the underlying problem within the network operations. Plenty of measures or KPIs will exhibit anomalies and an operator will still be overwhelmed with the analysis of the output of an anomaly detection engine. There should be a mechanism that would allow us to gauge the significance of an anomaly and would help indicate where the problem originated. However, we can receive help here using an ensemble function that provides an unbiased, multi-variate deviation of a vector X from its mean vector μ. X and μ in turn are composed of normalized Z-scores coming from various anomaly detection algorithms. As a result, the function provides a single score across many dimensions independently on the type of anomaly and the scale of a metric.

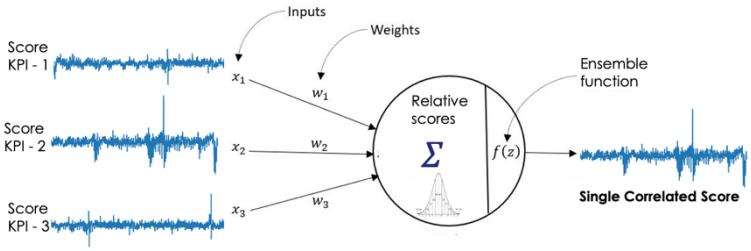

Figure 6. Ensemble function normalising and aggregating multiple anomalies across many dimensions.

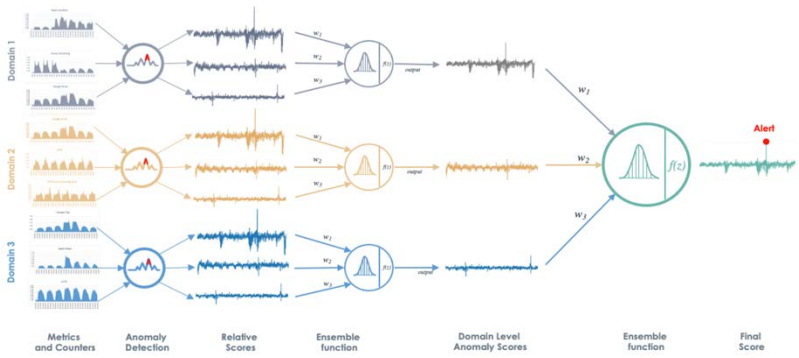

Figure 5. Hierarchical application of the ensemble function to create a single score across multiple domains.

An additional benefit of this approach is that these functions can be applied in a hierarchical way, providing a single score across multiple network domains. The hierarchy of domains and metrics need to be configured at the time of implementation and typically would be based on network topology or on different type of logical aggregations that are important from the network monitoring perspective. For example, a hierarchy can be constructed based on the type of operating system and the type of application being used or can be based on network topology or on a mixture of both approaches as presented in Figure 7.

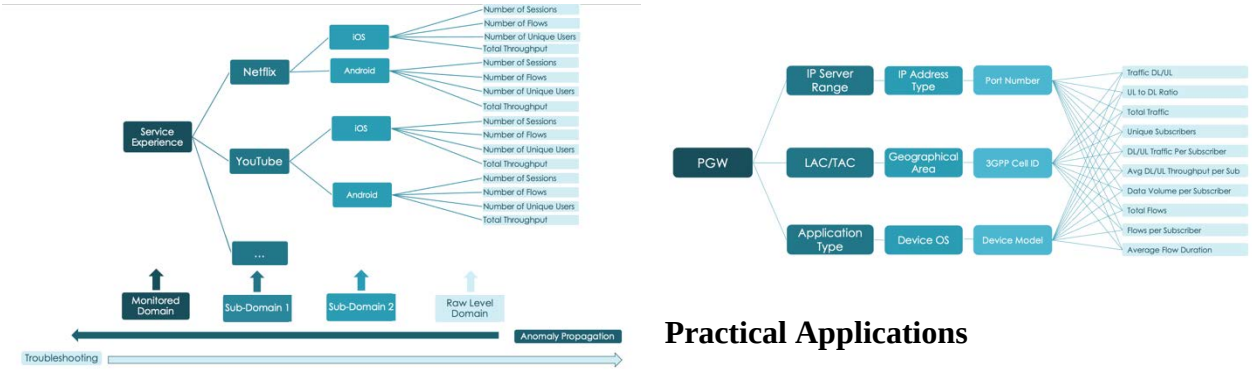

Practical Applications

Figure 7. Examples of different ways of defining domain hierarchies.

Detection of impacting BGP (Border Gateway Protocol) traffic switching

One of the applications of this approach that operators find useful involved that detection of service-impacting changes on the network. Figure 8 provides an example of an interconnect of numerous

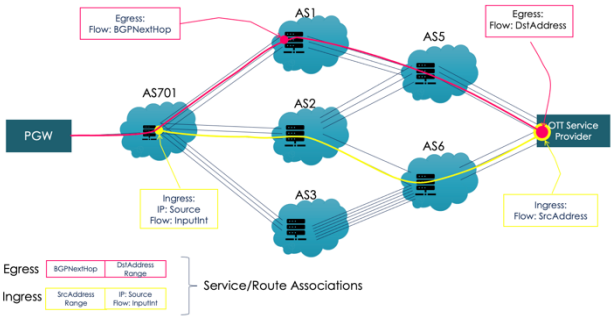

Figure 9. Detecting anomalies in BGP routing

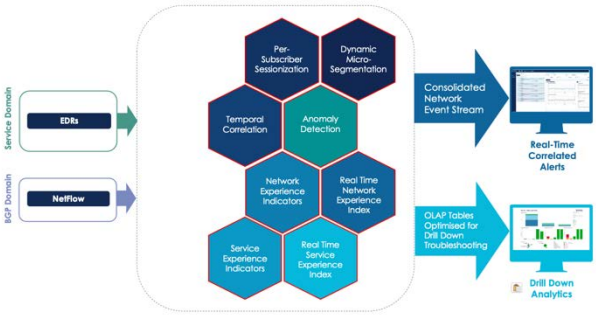

Figure 8. System diagram (OPS-IQ) for detection of service impacting network changes.

Autonomous Systems (AS) and ingress and egress traffic flow from a service or content provider. For the detection of service issues/anomalies, Event Data Records (EDRs) containing the user-plane information are ingested from a Packet Gateway (PGW), sessionized (i.e. all the data flows belonging to a single service are joined together) and aggregated on a service level with corresponding quality metrics. At the same time, Netflow data are ingested to monitor the traffic ingress on the ingress ports and traffic egress based on the BGPNextHop value. A system diagram for such solution is presented in Figure 9.

Such approach can also be used to warn of any potential threats associated with traffic routing changes.
In the scope covered by this paper, such changes in a private network would provide insights of real availability of communication nodes, which could trigger a response (e.g., enhance its capabilities or send someone to protect the asset), while on the scenario of interconnecting private bubbles via a public network could provide a warning of a malfunction of the public network or an anomalous increase of traffic inside of the bubbles.

Detection of Transmission Congestion or Performance issues.

Another example of application of such a system is to quickly discriminate capacity or performance issues resulting originating from RAN (Radio Access Network) or the transmission domain. In this example,

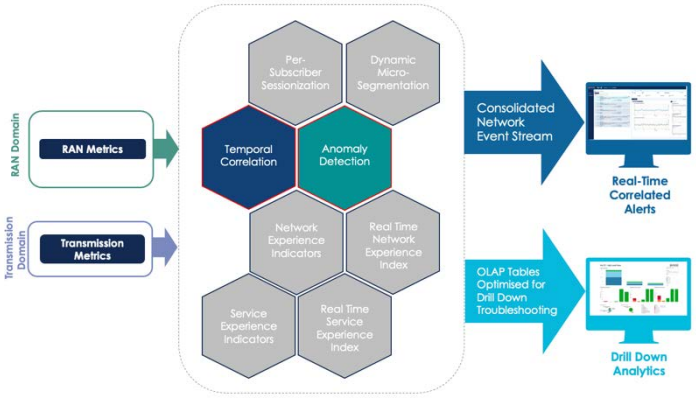

Figure 10. Thales platform (OPS-IQ) setup for Transport/Transmission use case.

performance metrics on a cell level coming from a RAN domain and the transmission metrics from a transport domain are fed into an anomaly detection and correlation engine. In the case where an anomaly were detected across a number of cells connected to the same transmission hub and at the same time an anomaly were detected in the transmission counters, an alert would be raised with the indication of specific network elements involved. In this case fewer platform modules are engaged as all the input data are in a timeseries format and hence no sessionisation of the data is required.

Any performance issue inside one of the bubbles, and even further traffic congestion, is always something to be avoided, as that can mean that critical information is not being delivered in due time. Data priority mechanisms can then be immediately activated to avoid the non-delivery of the most critical information, with several levels of quality of service.

Additionally, the detection of performance issues alongside an estimation or measurement of the traffic generated can be an early indicator of intrusion into the network or even information leakage.

Detection of DNS or TCP Signalling Tunnelling

Another use case involves detection of anomalies associated with a specific type of traffic protocol or application. For example, DNS traffic is characterised by large number of transactions but low volume of traffic. We have been able to detect scenarios where DNS traffic becomes abnormal due to DNS-tunnelling fraud where users hide the data payload within DNS packets that are always permitted and that are zero-rated, meaning they are not being charged for data transmission. We have also detected a similar type of fraud where the data payload is being tunnelled within TCP signalling packets (such as TCP SYN and SYN-ACK packets).

What's coming in 5G

The 5G SA network, that is currently being deployed in a commercial space around the world, brings revolutionary changes to the network architecture and to the way new services will be delivered. Network slicing brings a lot of benefits to private and military sectors, as it enables us to separate part of the network and dedicate it to a specific service, business or indeed to create a separated and secured military bubble. However, slices would still be managed by network providers and hence a visibility of slice performance and managing slice SLA (Service Level Agreement) is of vital importance.

The aspect of performance management in 5G has also been well thought out when defining standards by 3GPP. The 5GC (5G Core) network has been defined with Service Based Interface (SBI) at its heart, and what is more, the core network also received a new network element responsible for analytics called NWDAF (Network Data Analytics Function). NWDAF gathers events and statistics from each network function and from OAM (Operations, Administration and Maintenance) systems and provides analytical insights and predictions that can be directly used by the core network, application functions and even can be used for SLA management.

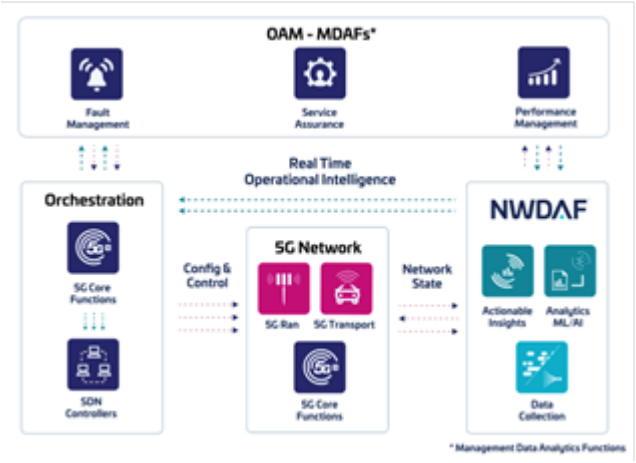

Figure 111. NWDAF in 5G Core Network

An interesting use case that can bring direct benefit in military installations is an outage notification system. NWDAF has a capability to track individual User Equipment (UE) and monitor its location and predict its moving trajectory. In case of an outage of a specific cell, NWDAF can provide in real time not only a list of users impacted by an incident, but also a list of users who most likely will arrive in the area in the near future. A description of such a use case is presented in Figure 12.

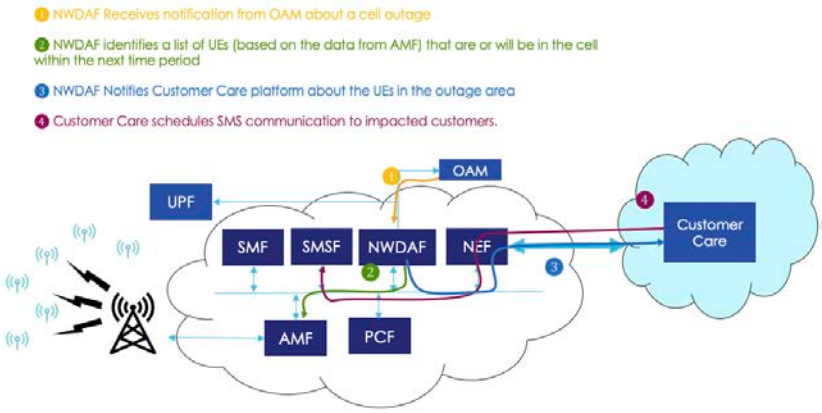

Figure 12. NWDAF Use case with automated outage notification system.

3. CONCLUSION

In the paper we have presented an innovative way of managing performance in complex mobile networks and described some of the algorithms that we have implemented in order to detect and cross-correlate anomalies across multiple network domains. We have also presented some use cases that could improve the operations in military operations. Eventually, we have also described what is coming in 5G networks and talked about 5G network slicing as an enhancement to the existing private LTE bubbles. At the end we discussed the arrival of a new network analytics function (NWDAF) into the 5G standards that will bring network automation to new level, as analytics output will be directly used by the network functions. NWDAF will also enable many analytics-based use cases and we have described how it can be used to automatically identify user equipment impacted by a radio cell outage.