VIAVI Solutions

White Paper

Mobile Networks: Automation for Optimized Performance

Mobile networks are becoming increasingly important worldwide as people transition to a more transient lifestyle. People now use mobile networks to work remotely, stream video, and access social media applications. Soon, mobile networks will play a major role in areas such as the Internet of Things (IoT), cloud computing, and vehicle communication.

This dependency on mobile networks has increased Quality of Experience (QoE) pressures on service providers at a time when bandwidth demands are also at an all-time high. How can service providers keep up with bandwidth needs and keep QoE at high levels?

Service providers are doing their best to meet these demands by making macro level adjustments to networks to achieve incremental improvements in performance. But this has come at a cost. Service providers are seeing profits decline as more money and staff are needed to keep networks running in this new, complex environment. Even with the increase in operating expenditures (OpEx), traditional network optimization is not enough to keep up with the dynamic nature of today's network traffic.

What is needed is a way to automate network performance to create major leaps in optimization on a granular level, while also decreasing OpEx and freeing up staff to maintain the infrastructure and plan for expanding the network to deliver greater capacities. Major advancements have been made in recent months to make automated optimization a reality. Let's take a closer look at the limitations of current network optimization methods, how automated optimization can overcome these limitations, and how this new method of optimizing networks can create a strategic advantage for service providers when the time comes to deploy 5G.

Challenges Facing Networks

As mobile networks continue to evolve, there are three main challenges that service providers face: interdependency, non-uniformity, and complexity. Each is a problem on its own, but together they create a network environment that is nearly impossible to optimize using traditional methods.

Many of the metrics used to optimize networks are now interdependent. Changing a parameter, or parameters, to enhance the characteristics in one part of the network can have implications on other characteristics in other parts of the network. For instance, trying to increase data throughput in a certain area could affect voice traffic – either positively or negatively – in the network.

This could also have a detrimental effect on design. Current designs that focus on one Key Performance Indicator (KPI) differ from designs that focus on other KPIs. This means that designs that focus on a specific KPI in isolation may or may not be the right choice for the overall performance of the network – especially as networks become increasingly non-uniform.

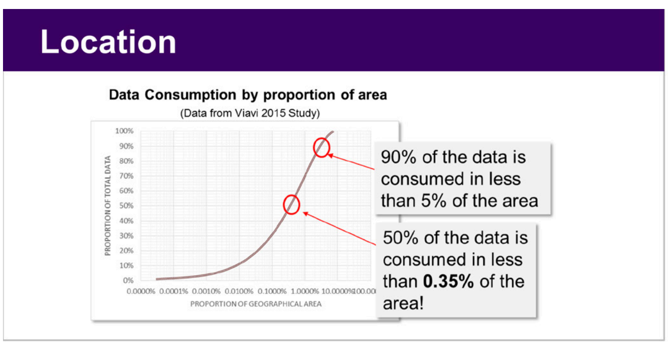

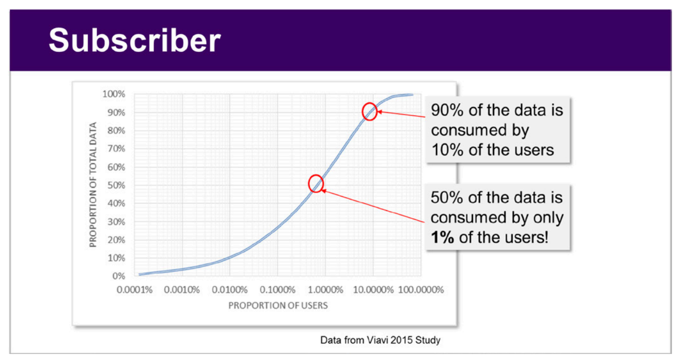

Extreme non-uniformity is the new normal for mobile networks as regular users become power users and the overall subscriber population becomes more mobile. According to the VIAVI Mobile Data Trends report, 50 percent of data is consumed by only one percent of users. In addition, 50 percent of data is consumed in less than one percent of a network area, and this area is constantly changing. This change can be dramatic. In extreme cases, the amount of data that a cell is expected to support can increase by orders of magnitude over a period of a few minutes.

This last data point is an important one. Not only has it become increasingly difficult to optimize networks because of non-uniformity, the non-uniformity is now dynamic. As this trend continues to grow, it will make it impossible to manually optimize networks in the future as this method cannot keep pace with the dynamic changes taking place.

This leads to the overall problem with optimizing mobile networks: complexity. Not only are subscribers using networks in new and dynamic ways, technologies such as LTE, VoLTE, and heterogeneous networks (HetNets) have added layers of complexity that mean that changes to a network layer will not only change how that layer responds to the traffic it must convey, but it will also change the way that layer interacts with other layers. For example, changing an LTE layer may make it more or less attractive at a given location to traffic on the 3G network, and vice-versa.

The number of tunable parameters is now enormous. For example, tuning just two parameters on each of 100 cells – where each parameter has 10 possible values – creates 10^{200} different ways these cells could be configured. That's more than the number of atoms in the observable universe!

Limitations of Network-Centric Optimization

The three main challenges put a spotlight on the limitations of current optimization methods. While networks have become increasingly complex and dynamic, most optimization efforts are still primarily network-centric: a problem is located using network statistics and then adjustments are made to the network parameters to solve the problem. This network-centric approach of characterizing a problem using network statistics and then making macro site parameter adjustments no longer works when optimization is needed on a more granular level. This approach is also less effective when the intention is to change the configuration such that the performance is improved, rather than solve a specific problem.

Taking this a step further, most macro-based adjustments create and maintain a baseline for overall network performance, but do little to optimize performance for specific locations within the network at any given time. For example, workers based in an office might tend to use voice services during the morning but then leave their office during lunch hour and go outside. While outside, their usage might migrate away from voice to data services. This illustrates the changing nature of the services demanded from the network and where they need to be delivered.

An effective optimization would have to configure the network to deliver an acceptable user experience for this cohort of users, not just during the work hours and lunch break, but also during the commute time, evenings, and weekends. At each of these times the usage profile will be different and the locations will generally change. Taking automation to the limit sees the network able to adapt its configuration as the day progresses in response to the changes in the demands placed on it.

But current optimization methods can only see macro locations based on overall network metrics. This creates "blind optimization" where multiple types of users at the various locations around the network are blended into one as the network tries to optimize an entire area. Doing so creates an imbalance where some users will have more resources than they need, while others will experience impaired usability.

Another limitation is the iterative approach toward optimization – making small, incremental changes over time – due to the inter-dependent nature of today's networks. This ensures that changes will not have an adverse effect on the network, but it also means that improvements are small with no major step changes in optimization. Most of these changes are use case driven and analyzed in isolation. If there is a problem with VoLTE performance for instance, current methods typically try to solve the problem in isolation without considering how it will affect other parameters such as data performance or energy consumption.

Drive testing is often used in network optimization. However, drive testing uses synthetic data and is OpEx heavy. It can also take a considerable amount of time and effort to come to a network design optimized for the drive test traffic rather than the commercial users of the network.

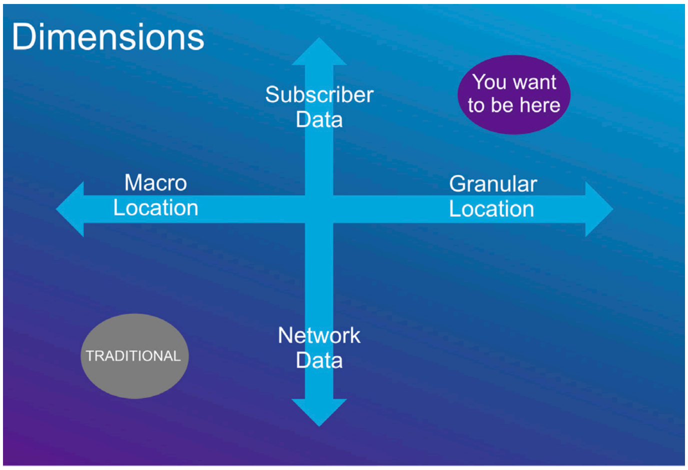

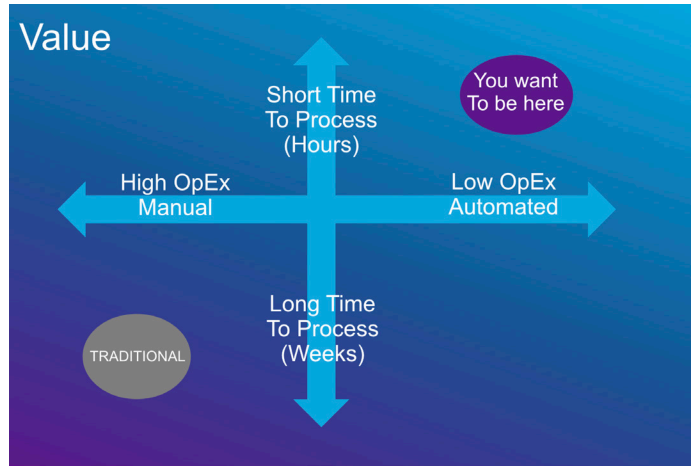

Most of all, today's network-centric methods only focus on the network itself and have limited ability to measure or enhance the subscriber experience of using a network.

Benefits of Automated Subscriber-Centric Optimization

New methods of optimization take the focus from the network to the subscriber. Subscriber-centric optimization considers where subscribers are located, how are they using the network, and what their current QoE is at any given time. But what must happen behind the scenes to make this happen?

Several advancements have made subscriber-centric optimization possible. Solutions can now collect, locate, store, and analyze data from mobile connection events, creating a repository of location intelligence from all subscribers throughout a network. This location intelligence is then transformed to deliver subscriber-centric performance engineering and Radio Access Network (RAN) planning information.

Most recently, subscriber-centric performance has been taken one step further by automating network performance optimization. This new automated subscriber-centric optimization addresses the network challenges created by interdependency, non-uniformity and complexity, and can keep up with increasingly dynamic traffic patterns.

Where traditional network optimization is a manual process and can take up to two weeks per site, automated optimization can optimize multiple sites at a time within hours rather than days. Where the focus of manual optimization must be a single site or a small group of sites, automated network optimization can focus on much larger clusters of hundreds of sites. Not only is the focus on larger clusters of sites possible with an automated approach, it is desirable since the exponential growth in possible parameterizations gives the optimization more scope to find configurations that maximize the performance for the mix of subscribers and applications in that region of the network. Once the area for optimization is selected, goals and success criteria are then established. KPI constraints and trade-off levels are then selected.

The optimization task is then scheduled – typically processing tens of millions of events based on subscriber data with granular location intelligence. If the results create the intended improvement, the changes can be actuated into the network. The result is a fast turnaround with major step improvements in optimization without adversely affecting other parts of the network.

Because this approach is automated, it also greatly reduces the staffing and OpEx needed to optimize a network.

Engineers are typically able to turn around optimized designs for large areas in a very short time. In addition, automated subscriber-centric optimization directly maps revenue to QoE to keep service providers profitable and subscribers happy.

In addition, the problems of interdependency and non-uniformity are overcome. Automated optimization can analyze KPIs in parallel and predict the impact of planned changes to make sure other parameters of the network will not be negatively affected. Algorithms calculate effects by predicting gains and the net costs of those gains to the network before any changes are made; and predictive decision making can resolve contradictions before they happen.

This more proactive approach saves time and prevents subscribers from experiencing negative events that are common using traditional, reactionary optimization methods. As an added benefit, the ability to use granular data at the subscriber level also allows network optimization to prioritize specific subscriber groups such as VIPs or high-net individuals.

In summary, traditional methods focus on network and synthetic data, are OpEx heavy, and take considerable effort and time to come to a conclusion that does not necessarily end up addressing the QoE and capacity issues. However, using subscriber-centric data ensures optimization is aligned with subscriber QoE, is OpEx light, and delivers network designs in a significantly shorter timeframe.

Automated Subscriber-Centric Optimization in Action

Automated optimization sounds good in theory, but does it work with real network traffic? Let's look at a few real-life examples.

A major mobile provider wanted to maximize data coverage and throughput by reducing the number of LTE data users on 3G. The goal was to improve data traffic volumes on an already optimized network while maintaining 3G voice service. The network had 233 cells across two Radio Network Controllers (RNCs).

Traditional

Collect	**Analyze**	**Actuate**	**Confirm**
Review areas Drive test prioritized sites Collect PM stats	Overshooters Drop Calls, Congestion, Load Balancing	Once manually correlated Then fix all sectors that have the issues	Re-drive problem areas Make final adjustments Check PM stats
TIME = 1 WEEK	**TIME = 3 WEEKS**	**TIME = 2 WEEKS**	**TIME = 1 WEEKS**

The Traditional Approach takes around 7 WEEKS with no subscriber data and limited granularity = HIGH OpEX & Effort

Automated Optimization

Collect	**Analyze**	**Actuate**	**Confirm**
Select Clusters Establish Goals & Constraints Schedule Task	Process Milllions of events Granular Subscriber Data Automatic analysis	Actuate optimization Design into the network	Compare results with predictions
TIME = 1 HOUR	**TIME = 1 HOUR**	**TIME = 1 HOUR**	**TIME = 1 HOUR**

The Automated Approach takes around 4 HOURS using granular subscriber data = Low OpEX & Effort

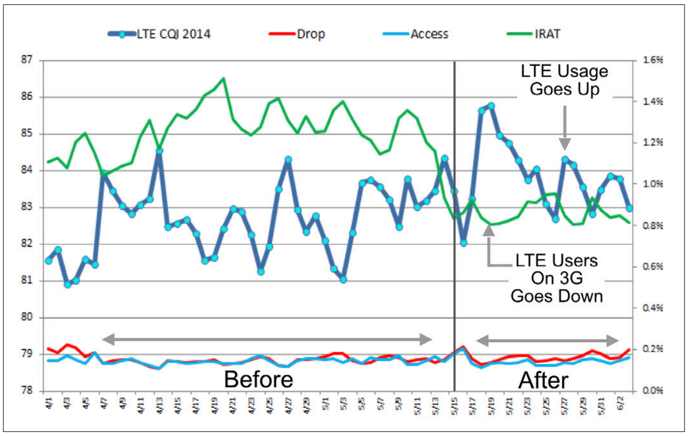

	Baseline	After
Accessibility	99.90%	99.92%
Accessibility (VoLTE)	99.82%	99.82%
Retainability	99.44%	99.46%
Retainability (VoLTE)	97.48%	98.03%
Mean throughput	6.56	7.46

Optimization targeted at VoLTE retainability

Automated optimization used subscriber-centric intelligence to analyze the current subscriber usage. Based on this intelligence, power changes were made to 67 cells, and 63 cells received antenna e-tilt changes. The result was a 1.3-point improvement in the LTE quality index and a 24 percent increase in data traffic volume – all without affecting 3G voice services. See diagram on left of page.

Another service provider wanted to maximize retainability of VoLTE calls and improve VoLTE throughput while maintaining accessibility. They also wanted to make sure the changes wouldn't impact data services.

Automated subscriber-centric optimization maintained VoLTE accessibility at 99.82 percent while improving VoLTE retainability from 97.48 percent to 98.03 percent. At the same time, the mean throughput improved by more than 13 percent. This was a major step change improvement without impacting data services. See diagram on right of page.

Voice and data are not the only uses of automated optimization. Service providers can also use it to optimize energy consumption to reduce OpEx without affecting subscriber services.

One service provider wanted to reduce energy consumption on their 3G network at major sites in a city while ensuring service availability. Automated optimization analyzed subscriber usage at key sites outside of normal hours and analyzed handset carrier capability. The solution also determined the optimal carrier configuration per site to optimize energy consumption while maintaining service levels. The result was a reduction in energy costs by 25 percent, saving the provider an estimated $2.4 million annually.

These step changes in optimization were all possible because real subscriber-centric intelligence was being used instead of traditional synthetic data. This allowed the service providers to see what the true results would be once the changes were actuated. Automated optimization allows engineers to establish specific goals to optimize aspects such as capacity, throughput, service drops or energy savings. Service providers can also focus on a select set of parameters for the most cost effective improvements such as only changing power or e-tilt parameters.

Automated Optimization and 5G

Subscriber-centric automation will become even more important as mobile networks become more complex. An analysis of a number of third-party industry resources shows that networks will see several major changes by 2025:

- 720 percent increase in video traffic
- 700 Billion things will be connected to the Internet
- 66 times increase in wireless traffic
- 2000 times increase in cloud objects
- 620 times increase in data analyzed in the cloud

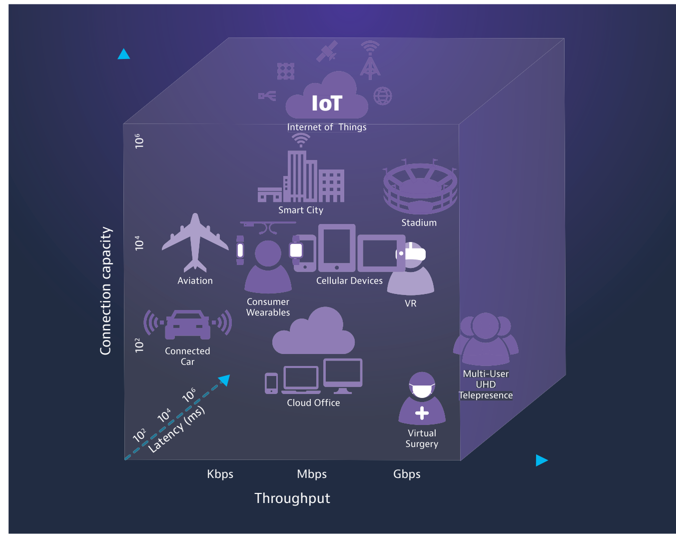

For mobile networks, service providers are looking to 5G to keep up with this changing demand. Although a lot of progress has been made, the standards for 5G have not been finalized. But the capabilities 5G must have to keep up with demand are staggering. According to the GSMA, 5G must accomplish:

- 1G to 10G connections to end points in the field
- Have 99.999 percent availability
- Reduce energy usage by 90 percent

A key characteristic of 5G is the expectation that it will be able to deliver connectivity to an even wider range of devices than are seen today. This will include public safety, and a plethora of IoT devices such as connected cars, smart meters and asset trackers. These devices will have a vast range of different requirements in terms of bandwidth, latency, jitter, reliability, and dynamics that will require a network to tailor the service to each set of subscribers and devices. The specific requirements for each group further compounds the problem of network-centric optimization as it's unable to discern the impact on each device and how it needs to change to meet QoE targets.

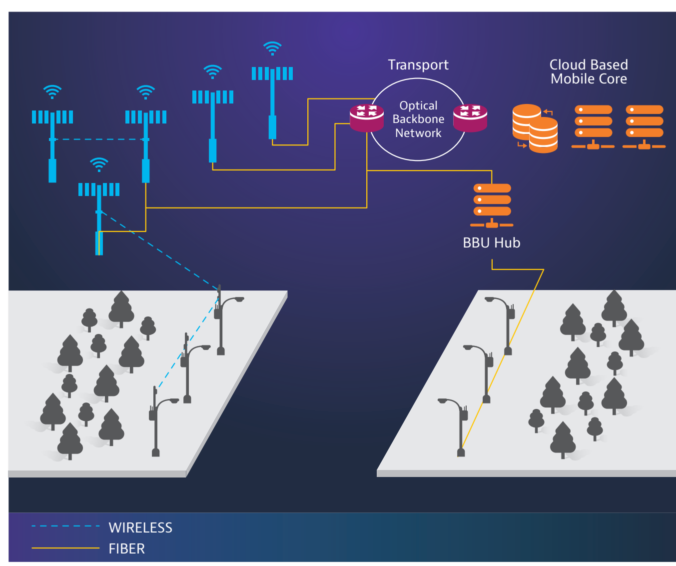

There will also be a trend towards RAN centralization and virtualization with the functionality of a traditional base station being split between centralized units and distributed units. In many cases these will need to be configured, managed and optimized in the context of their topology and transport constraints, and the subscribers they are serving. Advanced, coordinated radio transmission and reception schemes will be available which will provide better resilience to adverse radio conditions such as poor coverage and interference, but will come at a cost by placing more demands on the transport network.

The advent of 5G will also bring more use of Network Function Virtualization (NFV) and Software Defined Networks (SDN) to deliver network infrastructure. This will also require configuration, management and optimization. Other inflections such as Mobile Edge Computing will mean that functionality can be distributed and configured to meet constraints such as service latency and usage of transmission bandwidth.

5G will need to coexist and interwork with older technologies such as 2, 3 and 4G. Networks will gain another layer that must work optimally with the older technologies so that devices are still able to achieve their QoE targets. Any system that automates network optimization must perform effectively by taking advantage of all the layers, managing the selection of each layer, and transitions between them such that it sweats the assets and drives performance.

Taken together, these various developments make tomorrow's network more powerful by allowing devices more ways to achieve their various QoE needs. But this also creates a problem for management and optimization since there will be many more parameters to tune, the number of possible configurations explodes exponentially, and finding the optimal configurations becomes much harder.

The other impact of this increased configurability is the interdependency between different parts of the network. If changes are made in the RAN to address an interference problem, this may change the backhaul demands on a network. This issue is further compounded as some subscribers may derive service from different cells.

The relationship between a 5G network and the 2/3/4G layers may change as subscribers derive a service from these other layers in addition to – or instead of – the 5G layer. In addition, more devices may be attracted to the 5G layer. The network load could change as a result and place more demands on virtualized core elements.

Any optimization solution must be able to consider the holistic impact of configuration changes that are under consideration, as well as their ability to deliver the variety of QoE required by the different devices. Doing this effectively in the complex and configurable network will require advanced modelling of radio, RAN, transport and core elements along with mature configuration optimization capability to optimize the infrastructure and spectrum assets while delivering the required service.

The only way for this to happen is to automate optimization using subscriber-centric methods as a starting point and then add more automated features as they become available. Eventually, networks will need to have the capabilities of self-configuration, self-optimization and self-healing to keep up with subscriber demand and maintain a high level of QoE.

This may sound like science fiction, but it must happen and time is not on the industry's side. Currently, most service providers are planning mass deployments of 5G by 2020. Some service providers are already planning to make smaller deployments in 2018 and 2019. This means that automated subscriber-centric optimization is not a "nice to have" feature, but a vital step toward future networks. It's the only way service providers will be able to keep up with the complexity of networks and the dynamic traffic patterns of the future.

VIAVI Solutions

Contact Us **+1 844 GO VIAVI**
(+1 844 468 4284)

To reach the VIAVI office nearest you,
visit viavisolutions.com/contacts.

mobilenetworks-wp-maa-nse-ae
30186254 900 1017