VIAVI Solutions

White Paper

Harnessing Subscriber-Centric Optimization for the Next Generation of Self-Organizing Networks

Self-organizing networks (SON) have interested the cellular industry for many years. Work in this area started in 2006 at 3GPP with initial studies on the general concepts and requirements. There was a clear impetus to study SON—networks were becoming larger and more complex, and the maturing of 3G meant that operators had to grapple with the complexities of getting multiple radio access technologies to work and interact effectively. These heterogeneous networks had increasing numbers of configuration parameters, making their optimization very complex. The infrastructure was carrying increasing traffic, with a mix of services including voice, data, and SMS over 2G and 3G networks. This put pressure on CapEx to add capacity and OpEx to defer the need for capital spend by optimizing existing infrastructure.

In these early days, the overriding driver for SON was the competitive pressure to maintain the quality of experience (QoE) for an increasing subscriber base that depended more-and-more on devices for voice and data. Intense competition between operators meant that it was imperative to maintain QoE to avoid customer churn.

The case for SON is now even more compelling, and more challenging to deliver. Operators today face the challenge of rolling out and maturing an LTE layer while ensuring that their existing 2G and 3G investments work effectively in concert with the new layer.

Extreme Non-Uniformity in Today's Radio Access Networks (RAN)

The insatiable demand for data by subscribers is well known. This demand is even more challenging by virtue of the usage being massively non-uniform. For example, a study performed by VIAVI Solutions™ evaluated how data consumption was distributed throughout a network. The network was divided into 50 m^2 tiles and the total data usage by all subscribers in each tile summed. Figure 1 shows how demand for data is distributed between the different cells. Half of the data is consumed in only 0.35% of the geographical area of the network. This non-uniformity adds additional complexity to optimization. Extreme non-uniformity of demand means that site density will be similarly non-uniform. Often, an operator will resort to HetNet solutions with micro- and picocells and in-building solutions. This adds yet another set of challenges in the management and optimization of more network layers. The parameterization of this heterogeneous RAN serving a highly non-uniform and dynamic subscriber population makes the optimization challenge greater than ever before.

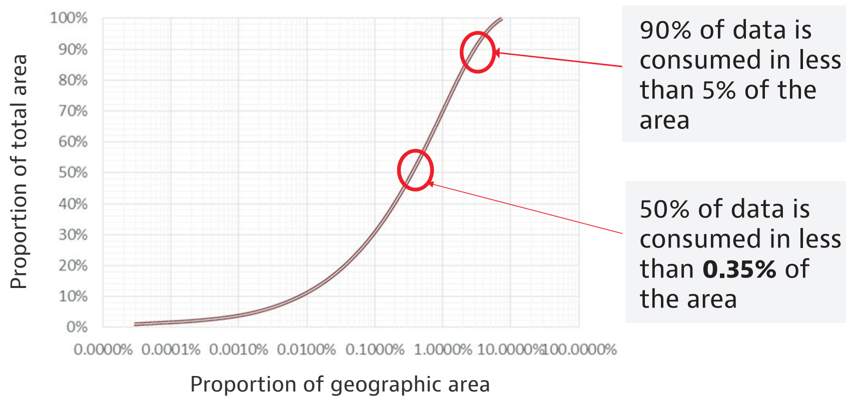

Figure 1. Non-uniformity in demand by location

Non-uniformity in demand by location is certainly a challenge. However, it is not the only source of non-uniformity with which a network operator must grapple. There is extreme non-uniformity in the temporal domain also, with demand for services fluctuating during the day and over the week, with changes often spanning orders of magnitude in a matter of a few minutes. This is on top of the relentless upward trend in demand that leads to the RAN requiring more or more capacity. There is also non-uniformity in the subscriber base; subscribers can be broken into cohorts, but each will have their own service usage profile, demand for resources, underlying QoS needs, costs, and value to the operator.

The Advent of RAN Splitting

Virtualization is taking place at all parts of the end-to-end network as the benefits of NFV and SDN drive their adoption. This virtualization extends to the RAN, with centralization of the RAN being a stepping stone towards the virtualization of the base station into a cloud RAN architecture. This evolution will give SON more scope to optimize the network, with new controls such as the functional split across the fronthaul and the use of coordinated transmission schemes.

The need to instrument these next-generation networks to gather information about conditions at numerous key points is driven by the requirements for mediation and orchestration to deal with fluctuations in demand, infrastructure failures, and new service provisioning. This couples with other innovations such as device-to-device communication that is expected to reduce latencies and lead to a more resilient network able to meet the needs of specialist applications such as emergency services. This is yet another aspect that will ultimately require optimization to achieve the best level of service for a given resource utilization. These disruptions in the industry mean that the distinction between SON in the RAN and optimization in the core, and even the data center, will start to break down and the industry will see a migration towards end-to-end SON.

The early advances at 3GPP have flourished into an ecosystem that sees a mix of NEMs and third parties providing a variety of optimization and SON solutions based on a wide variety of approaches. Each of these methods has strengths and weaknesses, but the changing nature of today's networks and the information available to understand them means that there is an opportunity to take the optimization of RAN to the next level. This paper explores these concepts.

Different Approaches to SON

There are many approaches to delivering SON to address the challenges operators face. Some of the distinctions are examined in this section.

Feedback-Based or Prediction-Based

One approach to SON is feedback-based. This approach is characterized by changes being made in the live network and the impact being assessed in how the performance of the live network changes. The configuration changes to be actuated in the network can be derived by a variety of means. They may be random or may be based on heuristic rules. The feedback-based approach is time consuming and error prone as it relies on stumbling upon better configurations by chance. This approach is in contrast to optimization based on prediction. This latter approach is characterized by the impact of predictions being assessed prior to any changes being actuated in the network. This method depends on having a reliable way to model the network and the impact of the changes made.

Optimization based on feedback is generally simpler than the prediction-based approach as the prediction introduces a step that is technically difficult to achieve accurately. This is in terms of some combination of the volumes of data required, the cycle time, or the algorithms that must be employed. However, since the impacts of changes are not generally known with much certainty until the changes are deployed to the live network, the optimization is unable to make substantial changes as these risk causing serious damage to the network design—with dramatic negative impacts on the subscriber experience.

Prediction-based optimization has a huge advantage over feedback-based optimization. This is because it predicts the impact of changes before they are deployed, and so it has a much greater scope to make more radical changes to the network design than the small changes that feedback-based approaches can make. It can make fundamental changes to the network design far more effectively than feedback-based methods.

Statistics-Based or Subscriber-Centric

Modern cellular RANs have a variety of telemetry available to the operator to facilitate network engineering, performance analysis, and optimization. An optimization solution will generally use some of this information, for example, as the basis of models for predicting the impact of changes or for assessing the impact in a feedback-based system. There are different types of these data. The first is aggregated statistical data such as averages and counts per cell, NodeB, and time period. Another, richer type of data is non-aggregated. It retains a high resolution of what each subscriber experiences throughout each voice call or data connection. This latter type is referred to as subscriber-centric data.

Subscriber-centric data is typically far more voluminous than statistical data and thus is more challenging to collect and manage. However, the high resolution is a major advantage that more than justifies the extra effort. Subscriber-centric data also benefits because it is geolocatable; the position of each subscriber can be estimated at any time the subscriber is active on the network. This is important when building optimization on predictive models, as the position of a subscriber allows the impact of parameter changes to be evaluated.

Building the most accurate predictive models for performance change is only possible using the highest resolution raw data, and that means that subscriber-centric data is a necessity.

Distributed/Centralized/Hybrid

Distributed SON (DSON) is characterized by the intelligence or algorithm execution residing at the network edge; in the eNodeB, for example. In contrast, centralized SON (CSON) sees the intelligence residing in the heart of the network.

DSON has an advantage because it minimizes the need to transfer data around the network. This generally reduces transport costs and cycle times. This approach can be suitable for various objectives, such as initial configuration of sites, configuration of physical cell identities (PCI), and neighbor list initialization. In contrast, CSON has an advantage because it can more reliably reconfigure aspects of the network that have less localized impacts. Changes to certain parameters in the network can have repercussions across a whole area of a network, sometimes including dozens of sites. For these optimization approaches, CSON is the best option.

As RANs evolve to 5G, capabilities such as mobile edge computing (MEC) will be standardized. This will open up the delivery of DSON applications to third parties beyond the NEMs. At this point, there will be a transition to a more hybrid SON (HSON). The evolved and virtualized RAN in 5G networks will be characterized by new parameters that SON can control, such as functional splits between remote and centralized units, the usage of coordinated transmission schemes, and the allocation of bandwidth to inter-node communication, fronthaul, and backhaul, for example. Instrumenting this will generate substantial amounts of data for driving optimization and thus a distributed component to the HSON will avoid the backhauling of large amounts of data to a centralized SON system.

Increasingly dynamic and heterogeneous networks are becoming characterized by parameter changes having secondary and tertiary side effects that are difficult to predict without sophisticated modeling. This means that a CSON component will complement the DSON components and create a potent hybrid optimization capability.

New Opportunities for the Next Generation of Optimization

Having explored the various approaches to RF optimization, we now focus on what the telecommunication industry needs and what technology is best suited for delivering it.

Focusing away from the Use Case and onto the Subscriber

Traditionally, SON optimization has focused on discrete use cases that address specific aspects of performance problems. For example, automatic neighbor relations (ANR) is concerned with achieving optimized neighbor lists. Neighbor lists should be complete so that every UE has a suitable handover target in every situation. At the same time, the lists should not be so long that the number of available slots is exhausted, or so the UE takes too long searching for the best handover target. Similarly, coverage and capacity optimization (CCO) is concerned with balancing what often must be traded off: the geographical coverage of the network and its capacity to carry voice and data for its users.

However, the radio interface of a RAN is highly non-linear and responds to stimuli or parameter changes in ways that are difficult, but not impossible, to predict accurately. Even small changes to network parameters can often have a multitude of effects. For example, an uptilt in an antenna or an increase in LTE reference signal power will attract more devices to take a service from that cell. The same change will also cause a change in the signal to noise ratio (SNR) for many subscribers. This change is not only for the subscribers that then take a service from that cell where, before the change, they would have been served by other cells. It will also change for the subscribers that do not change servers but are sharing the radio resource for their serving cell with a different set of subscribers.

This change in low-level RF characteristics for subscribers will change their experience of the services they use. Throughput may be affected for a variety of reasons. For example, throughput changes may be associated with the changes to the SNR described above but also may be because the level of cell congestion has changed, with cells becoming more or less congested.

A connection used for a VoLTE call may use different codecs with an associated change in the voice quality. Also affected will be the chance of admission failure and failure of established sessions. All these factors may impact customer satisfaction. Devices in that location may no longer be able to support VoLTE and may have to perform fallback to obtain a voice connection. Conversely, the VoLTE service may be available where previously it was not. The impact of the change will vary by subscriber depending on their radio conditions and what services they are accessing. The impact of the change will also vary by time of day with shifting traffic patterns as the day progresses.

This illustrates one of the key strengths of a subscriber-centric approach. When the network is understood to the level of the subscriber, subscribers can be geolocated and the resulting location used to make the calculated impact more accurate.

Figure 2 shows the complex potential impact on a subscriber resulting from a parameter change. What matters ultimately is the impact on the qualities of service and experience of each subscriber. This impact depends on the services used by that subscriber and the measures of service quality achieved for each of those services in terms of throughput, accessibility, and retainability. These are illustrated in green in the figure. These service measures will in turn be determined by the low-level radio including the signal strength and SNR which are in turn determined by the serving cell and the interference from users on other cells.

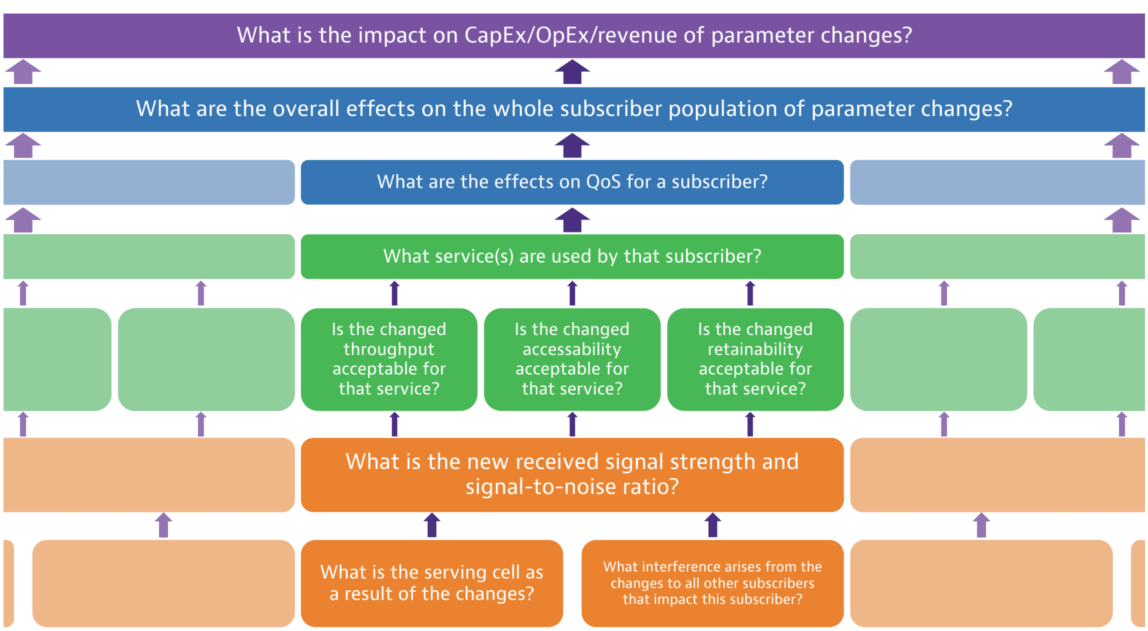

Figure 2. Potential impact on a subscriber resulting from a parameter change

This complexity will apply to all the subscribers in the network. Any parameter change can affect any part of this complex system, any of the subscribers, and ultimately the performance of the whole network, including those aspects of performance in which the operator is most interested. This means that each subscriber must be considered if the overall impact of changes is to be captured most accurately.

To illustrate the importance of this subscriber focus, Figures 3 and 4 show the impact of changing transmit powers and antenna tilts at three sites. The impact of these changes has been modeled for all the subscribers active in the network in a particular time interval. The figures show the impact on each subscriber for received signal strength and the burst rate, which is a proxy for the achievable throughput. Red indicates a degradation in that measure while green indicates an improvement. Gray indicates no change.

Looking at received signal strength, while there are some areas that are generally improved or generally degraded, there is a mix of improvement and degradation for both measures across the whole area. The modeled change here appears at first glance to be a degradation; the received signal strength is predicted to be worse for many subscribers shown in red. In contrast, for the same configuration, Figure 4 shows that the burst rate is generally increased for almost all subscribers. This is precisely because the signal strength is lower. The radio environment is cleaner with lower powers and so interference is less. Lower power to serve the subscribers means that there is more headroom available for conveying useful data. Lower interference means that there is more signal to noise available, enabling the use of modulation and coding schemes with higher throughput.

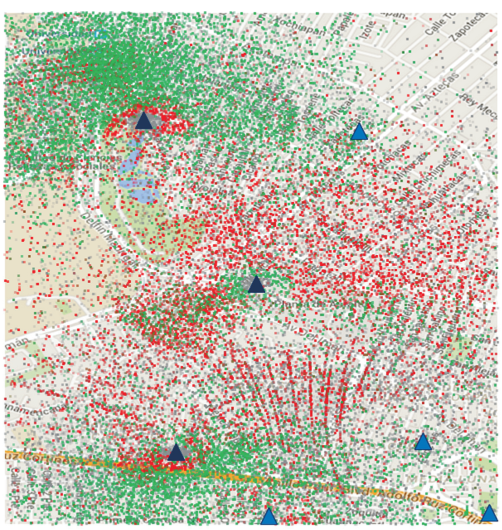

Figure 3. Impact of changes to signal strength

Figure 4. Impact of changes to burst rate

Many factors determine whether a subscriber experiences an improvement or degradation, including the strength of the serving cells and the interfering cells. These factors, in turn, are determined by the exact nature of the complex radio environment in that location. This granular structure in the response to changes is why building optimization on statistical representations of the network performance is so difficult to achieve with sufficient accuracy; these statistical models are simply unable to capture this granularity.

To have any hope of capturing the impact of a change to parameters means that this complexity must be recognized and thoroughly understood. Any attempt to do this without resorting to subscriber-centric approaches is not going to capture the full impact of changes. This complexity also tells us that the non-linearity in the network stands in the way of finding radically better network designs.

To significantly improve network performance, an operator must be prepared to deviate significantly from the current parameterization. That can be very risky to do in a feedback-based approach; radical re-parameterization can have substantial impacts on network performance, both positive and negative. To avoid the risk of substantial damage to network performance, one must use a predictions-based approach.

VIAVI has found that subscriber-centric approaches can make very accurate predictions about the impact of parameter changes. To illustrate this, a subscriber-centric optimization was performed on a 3G cluster and the cumulative distribution of burst rate was derived for the baseline subscriber population prior to any change. The same distribution was estimated for the network in the new state but before any change was made. The recommendations were then deployed to the network and the same distribution was measured in the optimized network.

Figure 5 shows that the burst rate improved after the change; in fact, throughput went up by 38% on average owing to a 1 dB overall average improvement in SNR. Even more interesting is the fact that the subscriber-centric predictive modeling was able to estimate accurately what the distribution of burst rate would be after the change—before that change was made. This is a key capability of the subscriber-centric approach. It makes the method extremely valuable for making changes that are more radical than possible using other approaches. It also enables changes that trade-off one measure against another in order to achieve objectives that are aligned to the business goals of the operator.

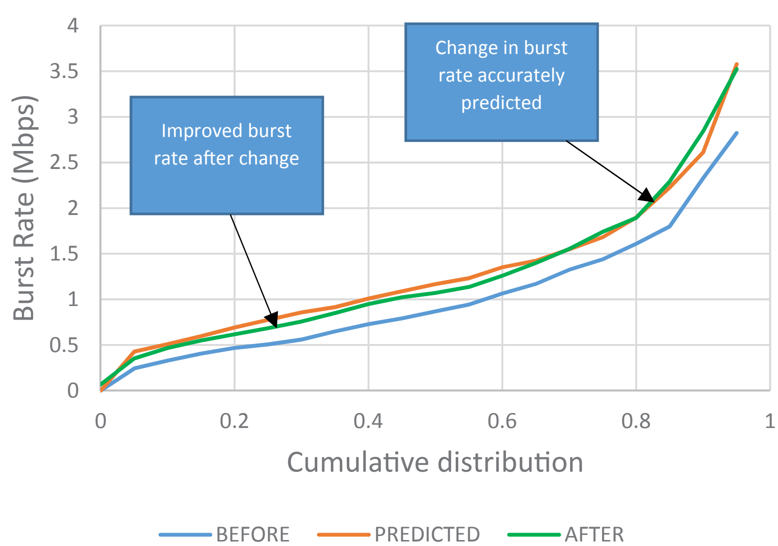

Figure 5. Prediction accuracy

Beyond the Silos of Use Cases

Every network operator is unique with their own customer profiles, technical challenges, and strategic goals. Even within operator groups, there may be distinct countries, markets, regions or areas, each with a set of challenges specific only to them by virtue of the terrain, morphology, demographics, engineering teams, commercial models, and priorities. A versatile SON solution must be highly adaptable so it can adjust to these different business realities. For example, one area may be suffering from a high level of drops and access failures while another ma be experiencing low data throughput or low voice capacity.

There can also be a disconnect between the real-world objectives of the operator and the traditional optimization and SON use cases that are used to address them. For example, dropped calls can have a variety of contributing factors such as poor coverage, excessive interference, missing neighbors, or misconfigured mobility parameters. To resolve this using traditional SON use cases would require deploying several use cases, and these would have to work in concert to find a solution. There is a danger that discrete silos of use case-based solutions may not work together; if neighbor lists are optimized for the current radio design, that radio design may be subsequently changed by the coverage and capacity optimization, rendering the neighbor list sub-optimal.

Overcoming this problem of silos requires an approach that does not focus on one aspect of the network performance at a time. Modeling the overall impact of parameter changes on the network, including the radio characteristics for each subscriber, the implications for mobility, and higher-level QoE measures, enables a more holistic assessment of the impact of proposed changes. This will achieve a higher-quality design. Using subscriber-centric optimization in conjunction with predictive modeling transcends silos and considers the overall impact of changes made.

Coexistence between SON Solutions

There are a variety of SON solutions on the market. For example, there are non-subscriber-centric solutions and many NEMs offer DSON solutions. DSON devolves the decision-making about optimization changes to the network edge, typically at the eNodeB. This is appropriate where the information required to drive the use case and changes made have only localized impact. However, changes can have far-reaching side effects on a complex RAN system. Changes with far-reaching side effects need a centralized agency to consider the full impact of the changes that are made and select changes that, implemented in concert, provide the highest possible performance.

Importantly, subscriber-centric optimization must co-exist with pre-existing solutions. It must be able to manage with changes being made independently and also must be able to add significant value alongside the pre-existing solution. For example, basing a SON solution on simulations of the characteristics of the networks such as models for signal propagation is difficult because the network will be in a state of flux outside of the complete control of the next-generation SON solution. Secondly, the next-generation SON solution must be able to deliver the real promise of SON: to select configuration changes that maximize benefits and surgically tune performance precisely to the business objectives of the network operator.

VIAVI has found that multiple RF optimization systems can coexist in the same network. Figure 6 shows how subscriber-centric optimization in the form of the VIAVI RF Shaping solution was applied to a 3G cluster of nearly 300 cells, allowing the pilot (CPICH) power to change in each cell across the cluster. Subscribers in the cluster were often unable to access HSDPA services. When they were able to access it, they typically achieved low data throughput rates. The network operator had to improve accessibility and throughput for the HSDPA layer. This was so important for the operator that in order to achieve these data objectives, the operator was willing to accept a small degradation in the voice drop rate.

There was an independent non-subscriber-centric SON solution active in that cluster at the time of this activity. Despite this, the subscriber-centric optimization was able to improve substantially the HSDPA throughput and the accessability by 16.4% and 7.6% respectively. The solid red line shows the date of the change.

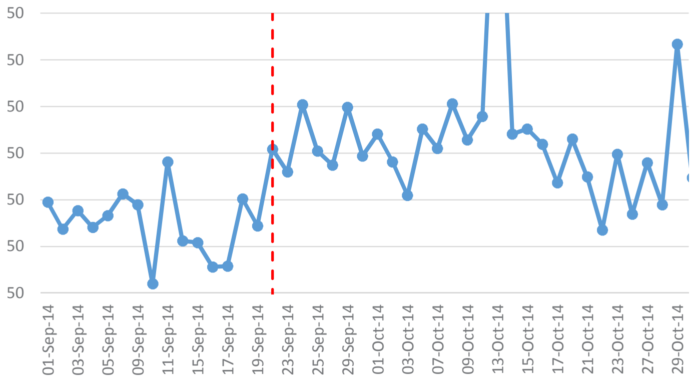

Figure 6. Improved throughput (16.4%) with VIAVI RF Shaping

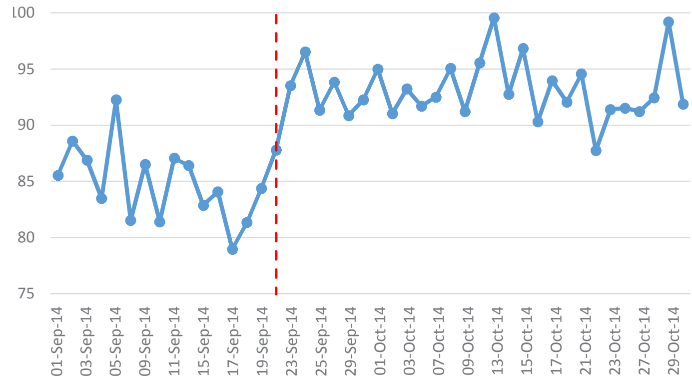

Figure 7. Improved accessibility (7.6%) with VIAVI RF Shaping

As directed by the customer, the optimization was allowed to degrade the voice call drop rate, and a very small degradation of 0.02 percentage points was seen after the change. However, the number of voice call Erlangs increased by 5.5% after the change, so this more than made up for the small increase in drop rate as the voice capacity was substantially improved. This was because the radio had been cleaned up for the HSDPA layer but the circuit-switched voice also benefitted from this improvement. These substantial increases in capacities and utilization of both the voice and the data layers translates directly into higher revenues and also has benefits for the operator in terms of deferred revenue.

Table 1. Performance improvements with VIAVI RF Shaping

HSDPA MAC throughput	+16.4%
% HSDPA accessibility	+7.6pp
HSDPA MAC data volume	+2%
% voice RAB drop rate	+0.02pp
% voice accessibility rate	+0.2pp
# CS voice Erlangs	+5.5%

Where there are existing non-subscriber-centric SON solutions in operation in a network, there are opportunities to step beyond these solutions to realize the full potential of performance improvements with a full subscriber-centric approach using predictive methods.

Taking Subscriber-Centric Methods to the Next Level

Once an operator looks beyond the silos of SON use cases, more exotic capabilities become a realistic possibility. Operator objectives sometimes necessitate trading off performance measures against each other. For example, an operator may want to increase capacity in exchange for some other measure. Or, there may be a radically wide-range of performance per subscriber, such as throughput. Some subscribers may have very high throughput, but it does not enhance their QoE. Meanwhile, their high throughput means that other subscribers have throughput that is so low that it impairs their customer experience. This is clearly not a desirable situation; it can be a valuable business goal to exchange the highest throughputs experienced by some users for acceptable throughput for the wider population.

This was the situation for a Tier 1 operator that VIAVI has been working with. A cluster of over 150 cells in their 3G network had a wide disparity between throughput achieved—many subscribers were not getting the quality of service they wanted. The operator needed to share the throughput more equitably to achieve a data service for more subscribers, along with improving the voice drop rate. Key to this was the willingness to allow a moderate degradation in peak throughput in order to maximize these improvements.

Table 2 shows the results. Data throughput rates were reduced as planned, but in exchange for this, a higher percentage of subscribers were able to access data services and the volume of data went up by nearly 9%. Not only was the data available to more subscribers, but there was a nearly 9% increase in the data capacity. This improvement for the data service was delivered by making a cleaner radio interface with well-defined coverage and lower interference. Because of this cleaner radio, the voice service also benefitted, achieving the desired voice drop rate improvement. The voice capacity was also improved and the Erlangs carried increased by over 4%, leading to more revenue for the operator.

Table 2. Data volume improvements with VIAVI RF Shaping

Voice drop rate	-0.16 percentage points
HSDPA accessibility	+0.55 percentage points
Voice accessibility	Unchanged
Voice Erlangs	+4.23%
HSDPA data throughput	-8.17%
HSDPA data volume	+8.87%

Because subscriber-centric optimization can accurately predict the impact of changes on all subscribers, it can choose a configuration that cleans the RF, trades off measures of performance against each other, and substantially increases the capacity.

As well as being able to trade-off performance measures to achieve operator business goals, subscriber-centric optimization is able to target particular services. For example, QoE for interactive and streaming services such as telephony and video services is highly vulnerable to poor performance in terms of throughput and congestion. In contrast, background services such as e-mail are not impacted as severely by degradation in these performance measures.

Discrimination between services was important for a Tier 1 VIAVI partner who has VoLTE enabled on their LTE network. The operator wanted to increase the retainability of VoLTE calls. Using subscriber-centric optimization methods on a cluster of nearly 260 sectors across 3 carriers, VIAVI was able to target the improvements specifically to the VoLTE traffic. As a result of this optimization, the drop call rate for VoLTE was reduced by over 20%. The overall accessibility and retainability for the network was maintained and specifically not degraded as a result of the targeted optimization.

Table 3. Targeting improvements for a specific service (VoLTE)

	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

Mean throughput was increased by 13.7% because the radio was cleaner after the design was deployed as shown in Figure 8. The cumulative distribution of the signal to noise as measured by the reference signal received quality (RSRQ) is shifted significantly to the right after the optimization, showing that the quality has increased.

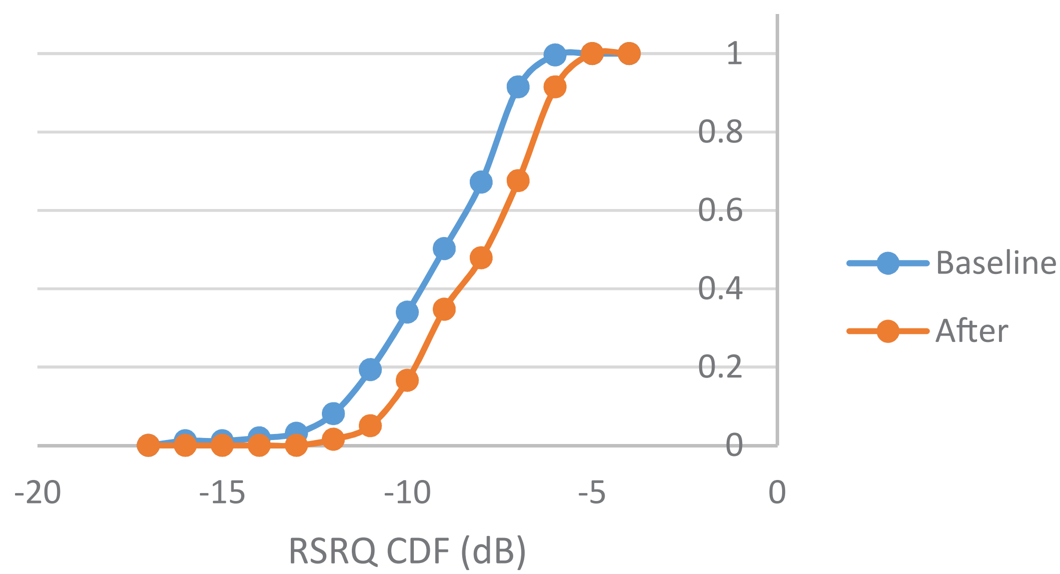

Figure 8. Mean throughput improvement

Conclusions

To realize the full potential of performance improvement, a SON solution must account for the complexity of the impacts of the changes it makes. This means taking into account the true impact on individual subscribers—this is imperative. Partitioning SON into narrowly-defined use cases such as CCO, ANR, and MLB is counter-productive as the side effects of each change will potentially be far-reaching. It is only this subscriber-centric mindset that can deal with the additional complexity of more exotic scenarios such as working with multiple generations of radio access technologies. It is here that more complex requirements such as spectrum refarming become a driver. These are requirements that subscriber-centric optimization addresses.

Making valid predictions about how a network will change in response to modifications of its state is the best approach to improving RAN performance. For the greatest impact, predictions must be based on a subscriber-centric approach. Only a focus on the subscriber can embrace the complexity, non-uniformity, and non-linearity of today's RAN. Moreover, a subscriber focus is able to take into account where handsets are located and accurately predict the impact of parameter changes. Subscriber-centric predictive modeling transcends traditional SON use cases and improves performance that, in turn, directly improves the subscriber experience in ways that subscribers care about—and that align with an operator's business goals.

The power of subscriber-centric predictive modeling goes further than providing stepwise improvements in network performance. Discriminating between different types of connections adds extra power to he technique. In this way, the method can favor particular service types, as in the VoLTE example shown previously. Other applications can target subscriber cohorts, such as roamers, corporate clients, or other high-value subscribers.

As operators grapple with the growing complexity of their multi-service, multi-carrier, highly heterogeneous infrastructure assets, subscriber-centric optimization is a valuable, key asset in the arsenal of network performance management solutions.

VIAVI Solutions

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

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

son-wp-nsd-nse-ae
30179568 000 1015