NOKIA

Machine Learning in 5G Networks for Enhanced Mobile Broadband and Industry Solutions

Insights Series

White paper

Achieving the full potential of 5G will depend on the ability to rapidly provision services and automatically adjust network and service parameters to offer the best user experience.

Machine Learning powers up this capability, allowing Communications Service Providers to offer new services to users, including industries with demanding low latency applications, while making the most efficient use of the Radio Access Network and cutting costs.

Contents

Executive Summary

5G networks are set to enable a whole new range of opportunities for Communications Service Providers (CSPs). As well as offering new services to their traditional markets, CSPs will be able to offer advanced services to new customers, with Machine Learning (ML) playing a fundamental part in their success.

Managing the complex, multilayered network characteristics of 5G will need advanced capabilities in both the Radio Access Network (RAN) and in management and orchestration. ML, with its ability to learn from previous operations, provides these capabilities. It helps to automate service provisioning for ultra-high bandwidth, Internet of Things (IoT) endpoints or ultra-low latency use cases.

From a business viewpoint, ML offers opportunities to reduce Operational Expenses (OPEX) through advanced automation. Advanced automation, in turn, is a means of providing further benefits such as better end-user experiences in a cost-effective manner, energy savings and higher spectral efficiency through advanced Multiple Input Multiple Output (MIMO) systems.

Figure 1. Nokia Future X architecture

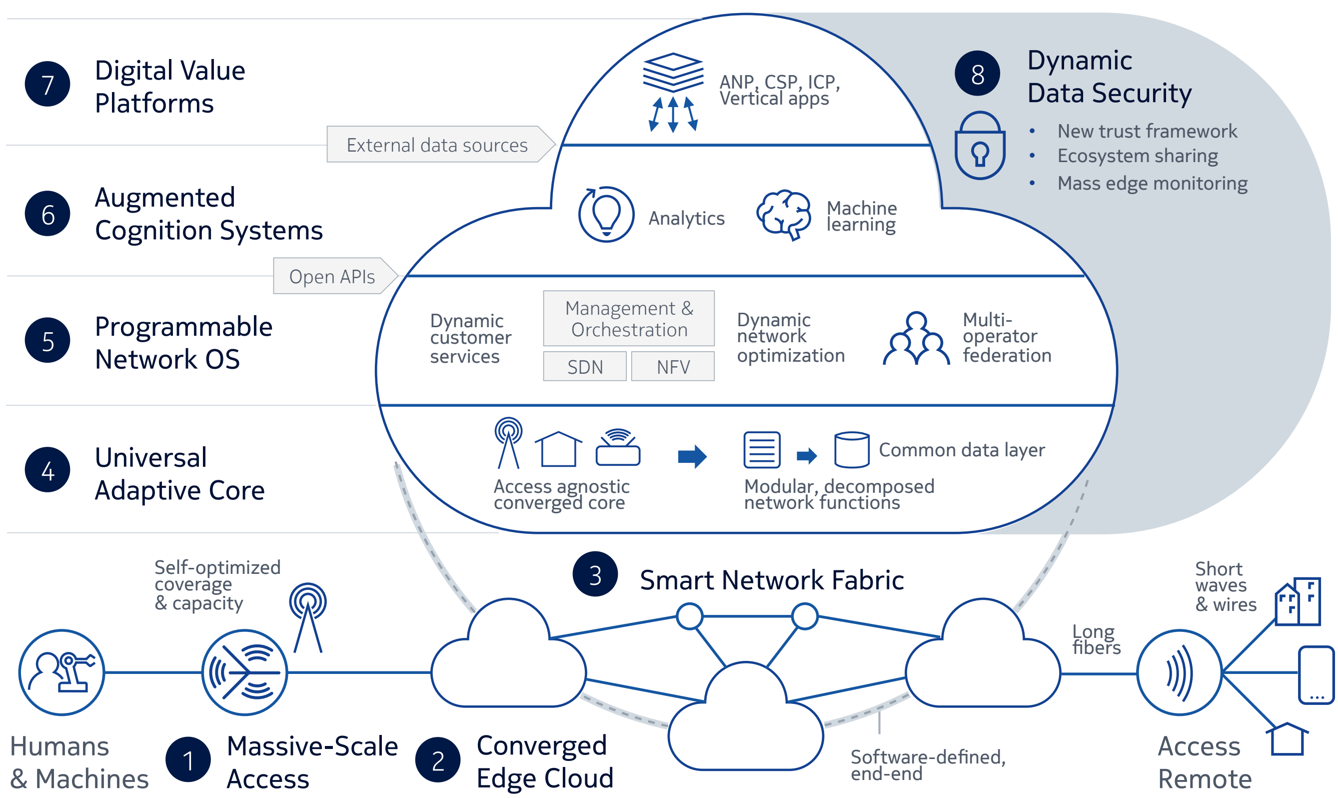

5G networks provide new capabilities for CSPs, helping them serve both current customers and new business. 5G supports a broad spectrum of service needs, ranging from ultra-high bandwidth, to Internet of Things (IoT) endpoints, to critical services requiring Ultra-Reliable Low-Latency Carrier (URLLC) bearers.

As well as enabling new communication services, 5G infrastructure supports new business opportunities for CSPs. These include hosting cloud applications at the network edge, for example, using radio channel state information to support completely new kinds of services.

The 5G network also brings greater technical complexity. Providing network slices efficiently requires a good knowledge of which resources will be available. There are many ways to deploy 5G RAN, including traditional base stations and cloud RAN. There are more options for optimization than in 4G due to the design of the 5G Layer 1 (L1) protocols, resulting in greater complexity.

The higher frequency bands used in 5G New Radio (NR) have different propagation characteristics to the existing frequency bands used for mobile broadband. For some optimization loops, latency requirements are strict (less than a millisecond). The execution environments for control algorithms range from next-generation Node B (gNB) to edge clouds and are mostly dictated by latency requirements. Consequently, designing control algorithms and continually adapting their parameters to specific cell environments requires significant resources. ML provides a way to automate this.

Applications built around ML models can replace some of the functionality previously done through software. Although running ML-based applications is more complex, with special requirements for support functionality, data collection and management, this is balanced by the enhanced capabilities they offer.

Machine Learning algorithms

For suitable use cases, ML can automatically identify the best response to a wide variety of inputs without explicitly being programmed. Indeed, ML algorithms learn by being trained on a range of different data inputs. Different types of ML algorithms can form parts of applications that need to react to input data: Unsupervised learning, supervised learning and Reinforcement Learning (RL). Deep Learning (DL) is a class of ML algorithms that can model complex inter-relationships between input features and target outputs for suitable problems.

Here, we look at using DL-based algorithms to build advanced automation in 5G networks. The different variants of DL, ranging from feed-forward networks, to autoencoders, to sequence models, to RL, to Generative Adversarial Networks (GAN), support a wide variety of use cases. DL can be used to meet the challenges of other types of algorithms - for example, in Deep Q learning, a neural network is used to make state space exploration more efficient.

The choice of DL algorithm depends on several factors - the control or management challenge, the latency requirements for control and the execution platform. For example, handling of complex input data with low latency may require DL because traditional algorithms might not be suitable. On the other hand, with less stringent delay requirements, ML might be used to augment and/or control conventional real-time algorithms or allow the identification and correction of problems before they occur.

Typically, an ML algorithm is used as a part of an application that uses the output of the model to implement a functionality. Use cases for ML applications include Radio Resource Management (RRM), anomaly detection and energy saving, as well as multiple applications of Radio Frequency (RF) fingerprinting. The Deep Q learning is a case in point. An algorithm using ML can be considered to have the ML model as one parameter to be managed, in addition to other parameters.

ML Framework

An ML algorithm alone cannot implement a use case in a mobile network, or any other real-world environment. A framework for managing the training and use of ML models is also crucial. In addition, an implementation of a real-world use case typically requires an application that uses the output of the ML model to achieve a goal. The use of DL as the heart of the application reduces the need for algorithms, since DL models encode relations between input features directly into the model, avoiding the need for traditional software.

Figure 2. Elements of an ML framework. Inference is performed within the endpoints. Training can be either centralized of distributed. Model sharing improves efficiency for training

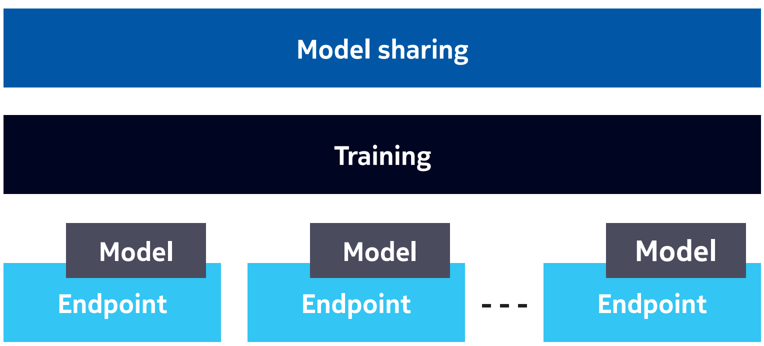

The availability of data is crucial for both inference (using the model) and training (creating the model). The requirements for data collection depend on the choice of algorithm, as well as the platform on which it is run. For example, for short latency applications, the data must be available when the application needs it. On the other hand, in high-performance parallelized training environments, ensuring an adequate volume of data is just as important as timing. Furthermore, a unified framework is needed to manage models in training and inference. Training of models can take place either online (continuous training) or offline (discrete training). The continuity of the inference model needs to be balanced with the processing needs of training.

Another consideration is the sharing of models in a distributed inference environment consisting of multiple endpoints. The simplest way of using ML is based on inference combined with offline training. In this approach, the model used in inference is trained (or re-trained) with data representative of real world use.

Models can be trained only on data relevant to a single endpoint or the data encountered by other endpoints. For the latter, transfer learning can be used to share models across sufficiently similar contexts. Transfer learning can be used to improve training performance, using a model pre-trained with either simulations or a sufficiently similar context. If the two contexts are similar enough, the inference model from one context may be used directly in another context, without training. Another option is federated learning, which combines distributed on-line learning with global combining of model updates.

On-line learning can be used to continuously update the inference model with input data. Ideally, federated learning would allow locally learned models to be shared with other inference contexts. The limited size of hard real-time inference models restricts how complex the model can be, limiting the use of federated learning in constrained environments.

Computer simulations can be used to aid model training. A coarse-grained simulation model can be used to train the lowest layers of a feed-forward network, corresponding to the major phenomena that a DL model could encounter in real life. A more advanced approach is a digital twin for a part of a network, which simulates advanced features such as propagation models and mobility patterns. Such a simulation can be used to give the DL model more detailed training.

ML and DL benefit from hardware acceleration. In inference, specialized edge inference accelerators allow the use of low-latency ML based control loops[1]. Nokia has developed ReefShark chipsets for the acceleration of ML in RAN. For offline training, general-purpose hardware accelerators such as parallelized Graphics Processing Units (GPUs) help train models with relatively large data volumes.

Full-stack ML

A consequence of adopting cloud platforms in 5G is that Network Management (NM) incorporates both NM and orchestration. For instance, setting up a network slice brings into play related resources, including the configuration of related Virtual Network Functions (VNFs). Conversely, network management needs to be aware of orchestration, which affects the virtual resources available. For both domains, the use of ML enhances the accuracy of predictions and thus the effectiveness of M&O. In addition to employing ML internally, M&O is used to configure and optimize parameters in ML enabled applications in the network.

In 5G we expect ML to be used at all levels, from an ML capable RAN, to network management automation, to orchestration. A multi-level ML capable system needs the ability to use advanced automation to solve specific use cases without jeopardizing the overall functionality. The key to achieving this is a deep understanding of the domain and focused problem solving with ML, combined with modelling and simulations of the system behavior. Nokia's end-to-end portfolio for 5G includes products and services for both networks and M&O layers, both with ML capabilities.

Finally, the use of ML affects many aspects of mobile networks, in addition to its role as an enabler for advanced automation. Network management needs to support the lifecycle of ML-enabled applications, including the management of ML models and parameterizing applications using ML. Professional services related to mobile networks need processes and competences related to ML, for example, for troubleshooting ML-based systems.

In this paper, we discuss the use of ML in the RAN and M&O. We then discuss ML aspects related to verticals as well as platform issues related to ML and conclude with a summary.

Radio Access Network (RAN)

Below, we discuss the use of ML in RAN from the view of the physical layer and RRM. We then discuss inference and training in RAN, as well as the importance of latency and conclude with selected use cases.

Physical layer

On the physical layer, the RAN faces specific challenges, such as symbol detection, channel estimation and digital predistortion, to name a few. These are all well-defined problems for which the desired outputs are known, although we may not have algorithms that achieve them. For example, for symbol detection during a training phase, we know the correct symbol decisions which should have been taken. This feature makes such problems straightforward from an ML point of view, as learning systems are typically trained to learn a mapping from the inputs to the corresponding outputs. For most narrow problems in communications, good models can be developed and the best algorithms derived. Hence, an ML-based approach must pass a very high bar for performance improvements.

However, ML can provide benefits other than performance improvements by reducing implementation complexity of certain algorithms and closing the gap between what should be implemented and what is done in products due to practical constraints - including compute resources and development time. Learned models are expected to achieve similar or better performance than the algorithms currently used, while possibly being implemented more efficiently in the form of neural networks with less overall computational complexity. Improvements of tens of percentage points in gain over non-ML cases are expected for selected use cases.

As an alternative to replacing well-established models, ML may be used to augment them, or assist in improving the determination of parameters. Joint optimization of isolated problems is not typically implemented because of the inherent complexity of algorithms. For example, channel estimation, equalization and signal detection are typically performed separately one after another. ML techniques may enable joint optimization, resulting in better overall performance.

Beyond these applications, ML enables new use cases by exploiting already available data for different purposes. As an example, the RF signal used for wireless communications is affected by moving persons and objects in the surrounding environment - this information is present in already available data such as the channel state information. A promising research area is the exploitation of channel state information to enable novel applications, such as user localization and detection of the presence of humans and their activities.

Independent of the type of problem, the real-time nature of most L1/L2 algorithms requires ML models, which operate and adapt at much higher speeds than those needed in other industries. Training techniques and neural network architectures designed for cloud-based environments cannot simply be used on devices with limited resources and hard time constraints. Thus, Nokia is actively developing hardware accelerators and solutions that allow for real-time inference and continuous training on embedded devices within the RAN and the edge cloud.

Although the full impact of ML on the RAN remains to be seen, Nokia believes that ML will play an increasingly important role and is investing heavily to be among the world leaders in this emerging area. In the long term, unified ML-based systems covering multiple layers in a system may replace larger portions of existing implementations.

Radio Resource Management

The task of (RRM) is, on the one hand, to maximize capacity through the best use of space diversity and the maximization of spectral efficiency and on the other, to meet service quality targets. The means of control include beamforming configuration and optimization, carrier selection, power control, inter-frequency load balancing and carrier aggregation, as well as optimization of the control and user layers.

RRM is associated with a broader set of tasks than the physical layer. Because of the complexity of challenges, heuristic techniques have been used to get the best solutions. Because it can adapt to data, ML offers a potential improvement over these heuristic approaches. However, from an ML point of view, these problems are "difficult" because one expects the ML model to essentially solve a complex optimization problem through trial-and-error. Putting such solutions into practice is challenging because learning on real systems tends to be slow, while today's simulation environments are unlikely to be sufficiently representative of the real world. One of the most critical questions is how we can transfer learning systems from simulations to the real world.

Now, let's look at some RAN-specific aspects of ML - inference and training and the latency requirements of control loops. They involve the timing requirements of control loops, as well as how ML capabilities relate to the source of data relevant to control. We refer to the components gNB (the 5G NodeB), Distributed Unit (DU) and Centralized Unit (CU), as well as real-time (RT) and non-real time (NRT) variants of RAN ML platforms. For example, depending on deployment scenario, an RT platform could be a far edge cloud.

Latency

Matching control loops that use ML to the infrastructure depends mostly on latency requirements. An overview is shown in Figure 3.

Figure 3. An overview of RAN platforms, corresponding time scales and relevant ML tasks

NRT platform	**O(s-h)**	**NRT inference**
RT platform	**O(10 ms)**	**RT inference**
CU	**O(ms)**	**Ultra RT inference**
DU	**O(sub ms)**	**Non-ML control**

For sub-millisecond latency levels, algorithms need to be implemented in the DU. On the other hand, the CU of the gNB can support millisecond level latencies. The respective hard real-time run-time environments are more limited than, for example, RIC or edge cloud, where models need to be computationally efficient. This limits the size of the model and hence also its generality. Consequently, the ML framework used to manage models for inference is important for gNB.

Latency targets of tens of milliseconds can be achieved on RAN RT server platforms, for example, a container-based edge cloud solution. It is a good match for inference, which requires non-hard RT performance.

Training and NRT can be performed in RAN NRT servers such as edge clouds or M&O platforms with relatively abundant processing power.

Use cases

A few example use cases of ML in RRM are described. The first is massive MIMO (mMIMO) scheduling, with the aim of optimizing spectral efficiency. An important aspect here is the computational cost of implementation. The challenge with mMIMO scheduling comes from the many combinations of input parameters, a result of the large numbers of antennas, beams and layers. A deep neural network can be used to learn a well-performing but complex algorithmic solution. A first prototype revealed a much lower implementation complexity, and hence latency, compared to a baseline implementation.

The second use case relates to beam mapping. The aim here is to maximize gain for User Equipment (UE) by making best use of beam patterns. One method is to predict mobility patterns, a task well suited to deep neural networks. Data on beam sequences, or the indices of the sequence of beams serving users, can be used to train a network to predict the best beam to serve the user. Neural Network (NN) techniques like that for Natural Language Processing (NLP) can be applied. Prediction of the next serving beam brings several benefits, including the reduction of control overheads and increases in spectral efficiency.

A third use cases relates to the localization of users. Channel state information available at a massive MIMO base station can be used with other sources of information to train neural networks to pinpoint users' positions in three dimensions. This positioning can be done with unprecedented accuracy, both indoors and outdoors as well as line-of-sight (LOS) and non-LOS scenarios. This approach can be used in stand-alone mode or to augment existing model-based solutions.

There are further RAN ML use cases in the optimization of user and control plane traffic, channel estimation, beamforming, detection of anomalies and the optimization of Quality of Experience (QoE).

Figure 4 shows an overview of RAN ML use cases in the context of 5G system architecture.

Figure 4. RAN ML use case overview in the context of RAN architecture

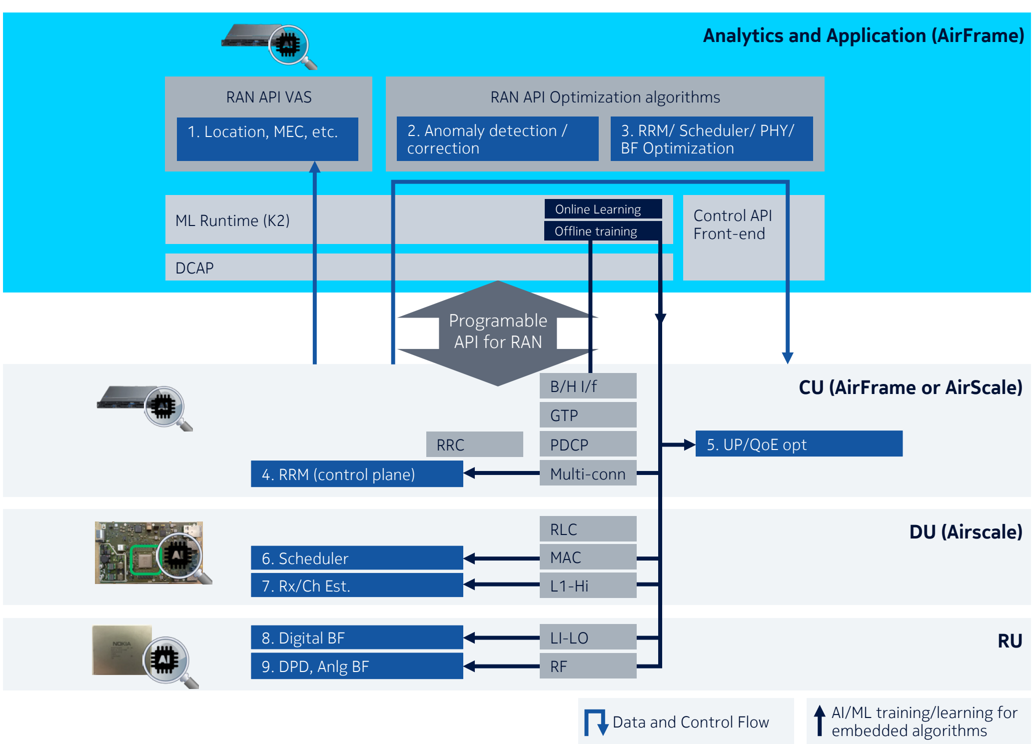

Management and Orchestration

The Management and Orchestration (M&O) domain is a complex area and of paramount importance to the success of 5G. M&O enables efficient use of the installed base while having the flexibility to support use cases and business models. The factors relevant to M&O in 5G include support for slice provisioning, increased technical complexity of networks, use of the flexibility provided by cloud and Software Defined Networking (SDN) platforms and new business models enabled by 5G. The provisioning of slices requires resources to be orchestrated and allocated end-to-end across numerous network domains.

Soon, the major types of AI applications in telecom will include[2]:

- Customer engagement
- Customer care
- Service operations
- Network optimization

For customer engagement, ML-enabled real-time analysis of customer related data is important and supports advanced services, such as proposing the right services at the right time. For customer care, ML aids the automation of customer interactions such as chatbots or NLP.

In service operations, ML can be used to automate the resolution of problems by identifying root causes and automatically triggering corrective actions.

For network optimization, ML can be used in Self-Organizing Networks (SONs), the automation of VNFs and complex analyses of scenarios for subscriber and data growth.

Next, we look at M&O from a wider perspective, focusing on the M&O of the 5G network. The next section investigates new business models of CSPs related to verticals.

Roles of humans and automation

A future M&O system will use the strengths of both humans and computers. Automated solutions offer consistency and the ability to react quickly to problems. For example, it is imperative to deal with problems affecting URLLC bearers promptly.

On the other hand, the automation of processes associated with a wide scope of information about the networks allows for proactive management to prevent problems occurring. The strengths of human users lie in their ability to take strategic decisions about the overall state of the network, for example, improving service quality in a specific area, or employing a more efficient energy saving approach in another.

Nokia's Connected Intelligence suite provides a unified AI approach for customer experience, services, networks and business. It provides centralized AI-augmented decisions, implemented with domain-specific mechanisms. The Connected Intelligence approach can be summarized as a centralized decision loop, which coordinates domain-specific optimization and execution. Nokia Bell Labs is studying novel approaches for Human-Computer Interface (HCI) for ML-capable M&O.

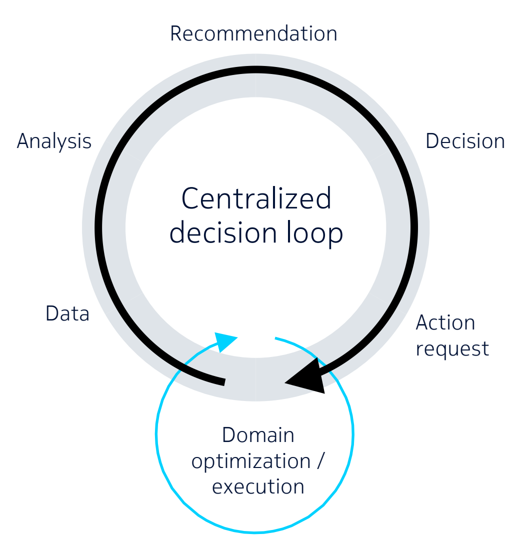

Figure 5. The principle of Connected Intelligence

Services

Network slices are enabled by M&O in 5G. They are associated with service quality parameters. Achieving efficient use of cloud resources is helped by predicting the availability of computing resources in cloud environments, a natural area of application for ML. The configuration and monitoring of resources used in a network slice employ the "M" aspects of M&O, that is, network management.

On the other hand, setting up network slices requires orchestration of virtual resources. In both cases, ML improves effectiveness through more accurate predictions. Such predictions can be used to plan slice provisioning, as well as in the management of network resources. Advanced monitoring is instrumental to more accurate predictability. Nokia Bell Labs is looking at topics relevant to network slices both in the areas of network management and orchestration. An example of this is application orchestration and rapid, flexible allocation of capacity across different types of 5G cloud platforms.

The goal for automated slice management has been defined in European Telecommunication Standard Institute's (ETSI's) Zero-Touch Service Management (ZSM). It covers the relevant domains, including cloud and SDN networks.

Automation

In 5G, network management needs to go beyond the static rule sets and policies of 4G SON functions. Machine Learning is again a natural choice for turning SON functions into Cognitive Functions (CFs). Nokia's EdenNet solution supports incorporation of ML into centralized SON functions. For developers, an ML capable Software Development Kit is available for developing SON functions. Use cases include prediction of network behavior, classification and clustering for data and anomaly detection. Nokia Bell Labs is developing methods for SON, including coordination and verification, as well as the use of transfer learning in M&O.

As with RRM, the implementation options of SON in 5G are determined by timescales. Tasks associated with real-time responses need to be implemented as Distributed SON (D-SON) functions, whereas for longer timescales, Centralized SON (C-SON) can be used. Optimization of slice related resources would be an example of a D-SON related use case, whereas Capacity and Coverage Optimization (CCO) can be handled with C-SON.

Automated learning can also be used more widely - the principle of self-learning allows past operational experiences, from both human and automated operations, to be used to aid learning. Overall, ML based automation allows transition from basic ZSM to true Augmented Intelligence, where automation enables flexible implementation of varying business scenarios while using 5G network resources effectively.

NOKIA

Industry use cases

The 5G Alliance for Connected Industries and Automation (5G ACAIA)[3] lists the following benefits of using 5G in industry:

1. Industry-grade Quality of Service (QoS) even under challenging industrial propagation conditions, with very rich multipath propagation and potentially significant interference.
2. The ability to deploy and operate industrial 5G networks in well-defined areas.
3. Standardized, open and flexible interfaces for seamless interoperability and handovers between public and industrial 5G networks.
4. End-to-end network slicing across heterogeneous technologies, countries and network operators, allowing dynamic and user-friendly establishment and release of application-specific network slices characterized by well-defined QoS and security properties.
5. Seamless integration into the existing connectivity infrastructure, taking the special characteristics and requirements of industrial applications into consideration.
6. Globally or regionally harmonized spectrum, for both licensed and license-exempt allocations.
7. Appropriate security concepts that consider both remote and local attacks and include device authentication and assurance of confidentiality, authentication and integrity for end-to-end messages.
8. A highly flexible and versatile air interface capable of satisfying the diverse requirements of the different use cases and applications, ranging from URLLC, through massive machine-type communication, to enhanced mobile broadband.
9. Support of multiple (well-separated) tenants using the same physical connectivity infrastructure in a factory.
10. The ability to monitor the current network state continuously in real time, even as a user, to take quick and automated action in the event of problems and to perform efficient root-cause analyses.
11. In-built indoor and outdoor user equipment localization with accuracy of at least 10 cm.

ML is a means of cost-effectively addressing the above issues. The use of 5G in RAN allows rapid adaptation to changing radio interface conditions and provides an advanced way of predicting resource use, an aid to resource management for network slices. The edge cloud infrastructure brings high-availability managed cloud infrastructure close to the radio access. Combined with low-latency / high-availability URLLC traffic, edge clouds can be used for hosting in-premises computations, for example, in industrial or traffic verticals. The edge cloud is also a natural platform for executing ML applications for verticals.

Figure 6. 5G cloud types

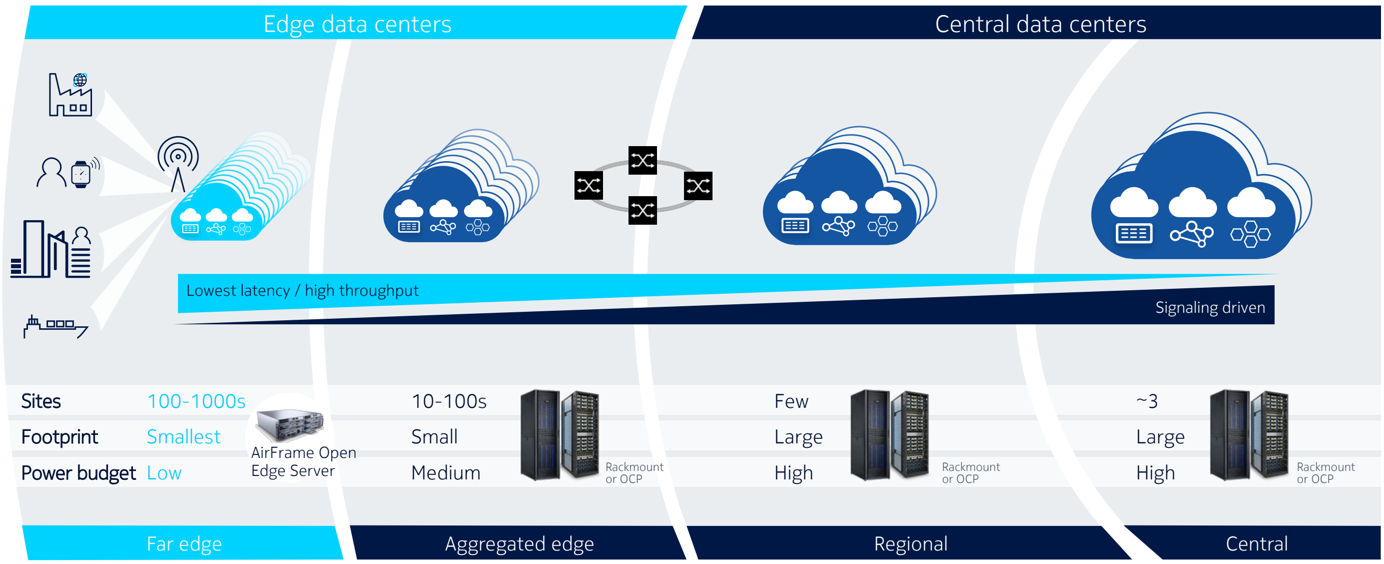

The 5G network provides multiple capabilities that either support new business models for CSPs or make implementing them easier. Edge cloud architecture allows the sharing of computing infrastructure between telecommunication functionalities and workloads of vertical applications. Multiplexing non-critical workloads with varying telecommunication loads can achieve a high utilization of infrastructure.

The URLLC traffic type is a critical enabler for many verticals such as factory floors. Multi-access Edge Computing (MEC) provides for low latencies, which are crucial for use cases such as real-time industrial remote control, especially where Augmented Reality (AR) is used. With Nokia's portfolio, MEC is realized in AirFrame Open Edge server platforms.

The IoT traffic type allows the deployment of many IoT endpoints in each unit area in an energy-efficient way, providing long battery life for endpoints.

For systems using the URLLC traffic type, changes in the environment may make it impossible to achieve the reliability targets. ML based learning techniques that combine inputs from sensors with network data can be used to predict potential communication failures, allowing corrective action before they occur.

Small Cell infrastructure operating on millimeter wave bands can be used to build dense base station topology. Combined with the ML-enabled processing of physical layer data, it can be used to provide insights into CSPs' customers' use cases. The relevant use cases may relate to drones, robots, vehicles, smart buildings, or smart cities, for example.

Summary and conclusions

Nokia provides an end-to-end portfolio of 5G products and services, using the capabilities of ML. Generally, the main benefit of using ML is the better accuracy of predictions compared to run-of-the-mill statistical models. In particular, the use of deep learning models allows the use of unified frameworks for training and using models in the production environment.

In the RAN, ML allows the implementation of closed loop control of broad parameter spaces over a wide range of timescales. This ability is crucial for the effective use of radio resources for demanding services such as URRLC in varying operational conditions. This is an important building block for effective use of a major feature of 5G networks, namely network slices. Slice provisioning benefits from the better predictability of resource use provided by ML, in addition to the adaptability of the RAN layer itself.

In addition to flexible service provisioning made possible by network slices, business agility benefits from ML in other ways in 5G. Small Cell infrastructures, industrial networks and innovations in the 5G physical layer open new business possibilities for connected industries, smart cities and other emerging areas of interest.

Further reading

Unleashing the potential of 5G: https://networks.nokia.com/5g

Nokia Future X: the ultra-optimized solution to de-risk your 5G deployment:
https://onestore.nokia.com/asset/205726

AirScale Radio Access: https://networks.nokia.com/products/airscale-radio-access

AirFrame data center solution: https://networks.nokia.com/solutions/airframe-data-center-solution

Edge cloud: https://networks.nokia.com/solutions/edge-cloud

Nokia AirFrame Cloud Infrastructure for Real-time applications (NCIR):
https://onestore.nokia.com/asset/205140

Networking solutions for industry - Go Allwhere: https://www.nokia.com/networks/go-allwhere.

Nokia ReefShark chipsets: https://networks.nokia.com/5g/reefshark

References

1. AI at the edge of innovation, https://www.nokia.com/blog/ai-edge-innovation/
2. The four most promising applications of Artificial Intelligence in telecom, https://www.nokia.com/blog/four-most-promising-applications-artificial-intelligence-telecom/.
3. 5G for connected industries and automation, 5G-ACAIA, https://www.5g-acia.org/index.php?id=5125

Abbreviations

5G ACAIA	5G Alliance for Connected Industries and Automation
AR	Augmented Reality
CF	Cognitive Function
C-RAN	Cloud RAN
C-SON	Centralized SON
CCO	Capacity and Coverage Optimization
D-SON	Distributed SON
CSP	Communication Service Provider
CU	Centralized Unit
DL	Deep Learning
DSP	Digital Service Provider
DU	Distributed Unit
ETSI	European Telecommunication Standardization Institute
gNB	next generation Node B
GPU	Graphics Processing Unit
HCI	Human/Computer Interface
IoT	Internet of Things
M&O	Management and Orchestration
MEC	Multi-access Edge Computing
ML	Machine Learning
NCIR	Nokia Cloud Infrastructure Real-time
NLP	Natural Language Processing
NM	Network Management
NN	Neural Network
NR	New Radio
NRT	Non-Real Time
OPEX	Operation Expenditure
OS	Operating System
RAN	Radio Access Network
RL	Reinforcement Learning
RRM	Radio Resource Management
RT	Real-Time
SDK	Software Development Kit
SDN	Software-Defined Networking
SON	Self-Organizing Networks
TRP	Transmission Reception Point
UPF	User Plane Function
URLLC	Ultra-Reliable Low Latency Carrier
VNF	Virtual Network Function
ZSM	(ETSI) Zero-touch Service Management

About Nokia

We create the technology to connect the world. Powered by the research and innovation of Nokia Bell Labs, we serve communications service providers, governments, large enterprises and consumers, with the industry's most complete, end-to-end portfolio of products, services and licensing.

From the enabling infrastructure for 5G and the Internet of Things, to emerging applications in digital health, we are shaping the future of technology to transform the human experience. networks.nokia.com

Nokia Oyj
Karaportti 3
FI-02610 Espoo, Finland
Tel. +358 (0) 10 44 88 000

Document code: SR1901032049EN (May) CID206223