Cambridge Consultants

INTEL® LINK PERFORMANCE PREDICTOR: DELIVERING BETTER SERVICE PERFORMANCE AND IMPROVED NETWORK EFFICIENCY WITH PREDICTIVE NETWORK AND SERVICE CHARACTERISATION

2019 REPORT PRODUCED ON BEHALF OF intel SK telecom

ABOUT THIS REPORT

Cambridge Consultants has developed an independent view of the potential use cases for Intel's Link Performance Predictor (LPP) technology.

This report outlines Cambridge Consultants' analysis and describes one of the first cases studies, detailing how SK Telecom is collaborating with Intel to train LPP to conduct automated network management.

By combining its deep technical expertise, broad telecoms experience and commercial insight, Cambridge Consultants provides a unique evaluation of LPP and its potential role in the emerging 5G ecosystem.

EXECUTIVE SUMMARY

- Expectations placed on emerging 5G technologies are high. To make the most of this potential, mobile operators will need to transition from a single purpose consumer broadband network to a versatile network supporting services as diverse as high-density sensor networks to real-time safety and mission-critical industrial applications.

- To justify the additional investment required to deploy 5G, operators will need to find ways to monetise higher value use cases. Many of these use cases will in turn have higher demands on reliability and performance than the traditional best efforts approach that has prevailed for consumer mobile broadband services.

- Technologies that integrate network layer and service layer performance in real time have the potential to deliver a differentiated level of service which is valuable in many use cases.

- Video streaming services account for most of today's mobile network traffic and this is expected to continue to be the case for the foreseeable future. This traffic is often tolerant to some delay, within buffering limits, leaving room for operators to shift this demand in time or space. Some other types of traffic, such as bulk software downloads, whilst a smaller proportion of overall traffic, are even less sensitive to time shifts.

- Where operators can smooth demand on network capacity by nudging services towards more efficient use of the network resources, peak infrastructure requirements are reduced and network utilisation is improved. Operators can use this advantage to either differentiate in quality of service or reduce network costs.

- Intel is addressing these opportunities with its Link Performance Predictor (LPP) technology. LPP uses cutting-edge machine learning techniques and a rich set of historical and real-time data feeds to dynamically predict the quality of any given RAN link and optimise application level behaviour. One of the first uses of the technology has been with SK Telecom, which is collaborating with Intel to train LPP to conduct automated network management.

- The LPP prediction engine can also be used for other network planning and management tasks to further improve network performance and reduce operating costs.

ADVANCED NETWORKING TECHNOLOGIES PROVIDE GREAT PROMISE

In 2019, we expect to see the first commercial deployments of Non-Standalone 5G in markets across Asia, North America and Europe. With the Standalone 5G New Radio specifications agreed in 2018 we are also edging ever closer to the 5G vision of perceived limitless connectivity.

Expectations are enormous. The ITU's specification for 5G sets out an order of magnitude improvement in data rates, device density, latency and power consumption. Many in the telecom value chain are keen to start exploiting the possibilities this offers for improved performance and fundamentally new services, such as:

- Gigabit wireless broadband to the home
- Smart factories powered by the Industrial Internet of Things
- Rich experiential media through Virtual and Augmented Reality
- Safe, efficient autonomous vehicles
- Cloud gaming

The ongoing transition to 4G has enabled, and simultaneously fuelled, an acceleration in demand for capacity, particularly video content. Building additional infrastructure and reacting to variations in network conditions (performance and congestion) as they occur in a fairly linear way has often been enough to respond to this adequately.

This was technologically and commercially challenging. Nevertheless, to exploit the promises often associated with 5G, operators will now need to undertake a far larger transition from this 'single purpose' broadband network to a versatile network supporting services as diverse as high-density sensor networks to real-time safety and mission-critical industrial applications.

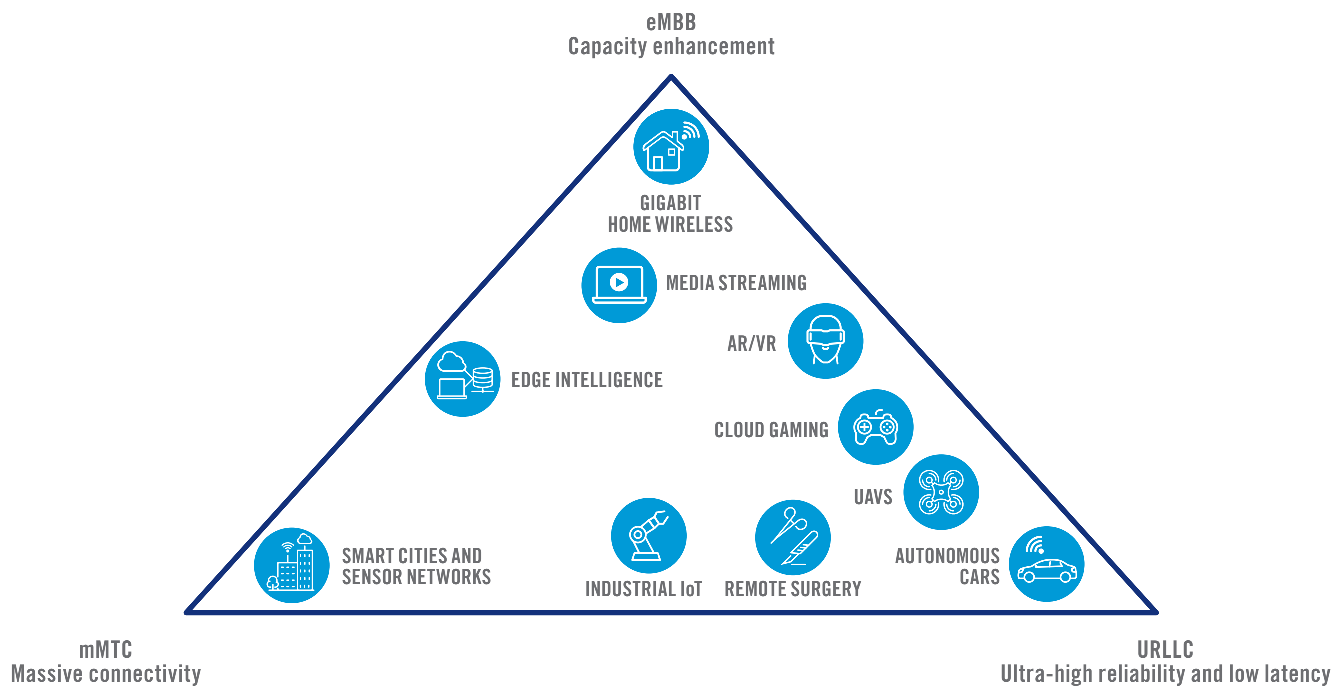

5G is characterised by a step change in performance

THE FUNDAMENTALS OF THE TELECOM VALUE CHAIN REMAIN CHALLENGING

Whilst the vision for 5G has generated significant enthusiasm, even beyond the telecom industry, the fundamentals of the telecom value chain remain challenging.

Peak performance has increased faster than average speeds, which have stalled at around 45Mbps[1] and whilst coverage is much improved, blackspots remain. The investment required to continue improving capacity and coverage is hard to justify for existing use cases. Many of the consumer broadband uses cases are relatively low value, as the declining Average Revenue Per User (ARPU) suggests, and this dominates network utilisation.

Higher-value services, such as mission-critical industrial or commercial applications, have been difficult for telecom operators to access, since existing networks operate largely on a best-efforts basis. Service performance and network efficiency both operate below optimal levels as services are unable to exploit optimum network performance – and this is because services operate without any awareness of the underlying network capabilities or conditions. Even when network information is exposed to the service layers, this is typically presented as static information rather than a dynamic, or even forward-looking, view.

With growth in demand continuing to accelerate (as shown in the chart), driven by the success of Over The Top (OTT) services, there is limited incentive to design these services with much regard to network efficiency. Consequently, operators face significant challenges to their business model ahead. If this traditional best-efforts approach continues, it is likely that services will continue to be exposed to variations in performance ultimately limiting their value to subscribers. Furthermore, services will continue to place increasing demands on network performance without necessarily resulting in increasing operator revenues.

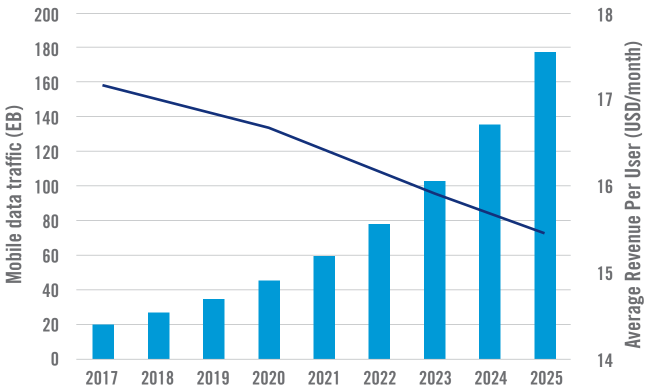

Consumers continue to demand more for less
SOURCE: GSMA Intelligence, Ericsson

Application developers and network operators could both deliver better services, with more efficient use of the available network resources, if services have access to the predicted network performance over time. Breaking down traditional barriers between services and the networks that support them will allow service performance to be improved in ways that were previously not possible. This will lay the path to more efficient network utilisation.

1 OpenSignal, State of LTE 2018, https://opensignal.com/reports/2018/02/state-of-lte

INTEL'S LINK PERFORMANCE PREDICTOR (LPP) PROMISES TO IMPROVE THE QUALITY AND EFFICIENCY OF SERVICE DELIVERY OVER MOBILE ACCESS NETWORKS

MOBILE ACCESS NETWORKS PLAY AN INCREASINGLY IMPORTANT ROLE FOR DIGITAL SERVICE DELIVERY. MEANWHILE PEAK PERFORMANCE IS IMPROVING FASTER THAN AVERAGE PERFORMANCE.

According to Cisco's VNI global forecast, global mobile traffic accounted for 9% of total IP traffic in 2017 and is expected to grow to 20% of total IP traffic by 2022. Operators continue to invest in additional network capacity to keep up with this rising demand.

Meanwhile, the capabilities of the wireless technologies underpinning mobile access networks also continue to improve rapidly. However, in a similar way to fixed access networks, the realised performance in mobile networks depends on the link quality. In mobile networks the variation in realised performance is more pronounced as the link quality conditions are more dynamic, particularly when the location of the mobile terminal is not static. Many of the developments that enable the performance improvements in mobile access networks rely on careful exploitation of these variations in channel conditions as they occur. As a result, the performance realised over mobile networks is increasingly dependent on the conditions experienced on the wireless link.

Today's services have typically either ignored this variation or relied on reactive adaptations. For example, MPEG's Dynamic Adaptive Streaming over HTTP (MPEG-DASH) and Apple's HTTP Live Streaming (HLS) adapt to changing network conditions to provide good quality playback with fewer stalls or re-buffering events.

LPP USES A RICH SET OF HISTORICAL AND REAL-TIME DATA FEEDS TO BUILD A VIEW OF THE NETWORK PERFORMANCE THAT SERVICES WILL EXPERIENCE OVER TIME

An ideal network delivers peak performance at all times. Failing that, the next best thing is to accurately predict how network performance will change over time. In many cases, accurate predictions are difficult to achieve with a high degree of confidence. However, in many instances, through a combination of historical and real-time data feeds, mobile network link quality can be predicted to a level of confidence and accuracy to guide service behaviour, such as content buffering and background data transfers. To this end, LPP uses a rich set of historical and real-time data feeds:

- **Historical data:** this is the performance experienced in a particular location at specific times on specific days and under specific conditions.

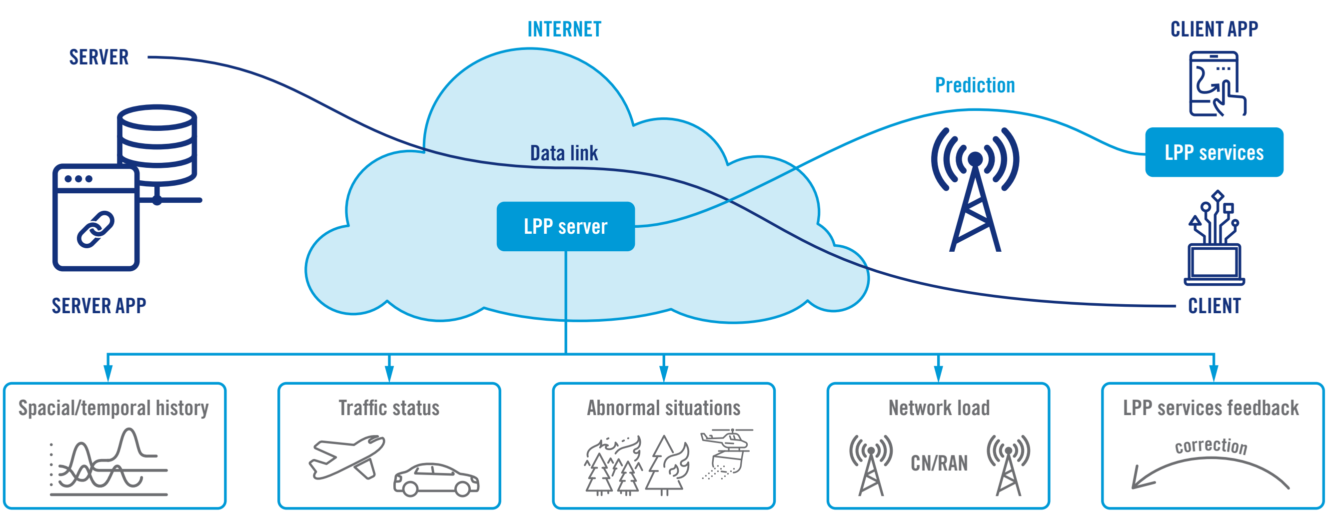

- **Real-time data:** this includes access to information such as live traffic status, abnormal events, network loading, radio conditions, device location, device motion and routing information (if known).

LPP uses these data feeds collected across a wide range of mobile devices, allowing LPP to build a view of the network performance that services will experience over time. In its most basic form the result is conceptually similar to a smartphone real-time navigation app. While the latter provides up-to-date information about current and forecasted traffic conditions, LPP offers information about current and forecasted link quality. Services can then subscribe to predictions that have direct relevance to the service performance. The LPP server responds to the client and/or server, indicating any significant changes in performance over time.

USING INTEL'S LPP PREDICTED NETWORK PERFORMANCE TO DRIVE MORE EFFICIENT USE OF MOBILE ACCESS NETWORK RESOURCES

By providing access to the predicted network performance offered by Intel's LPP, mobile operators can work in collaboration with key content providers to encourage more efficient use of mobile access network resources. In addition to traffic generated in response to user interactions, many applications generate regular background traffic even when there is no user interaction.

LPP enables mobile network efficiency gains by managing this background traffic. In this scenario, LPP's goal is to save network resources by scheduling this background traffic to coincide with peaks in realised network performance. This could, for example, correspond to mobile device proximity to a serving base station or a period during which the mobile network has sufficient capacity headroom.

IMPROVING QUALITY OF SERVICE OVER MOBILE NETWORKS USING INTEL® LPP

Quality of service is traditionally associated with network performance metrics. Like speed limits on roads, static network performance metrics only provide a partial, often incorrect, and at best indicative, view of the experienced quality. Following the Google Traffic analogy, by gaining access to current and predicted network performance, drivers (or their satnav route planner) can adjust their route to optimise their experience. Likewise, many services could adjust their mode of operation if they have advance warning regarding any significant changes in expected network performance.

LPP introduces a mediation layer between services and the underlying mobile access network, providing a simplified, abstracted view of the mobile access network link quality. This in turn allows services such as streaming content, real-time services and route optimisation to make proactive decisions to improve quality of service.

LPP for streaming content: using LPP, streaming buffers can adapt to cover for poor coverage areas that will likely be encountered en-route. Content providers and mobile operators can also couple this with Content Distribution Network (CDN) optimisation. With the help of LPP, required content can be pre-fetched to an MEC node closer to the user's anticipated point of consumption.

LPP for real time services: mobile operators and OTT real-time service providers can make use of LPP to provide advance warning on when, for example, a voice connection will be lost, and how long the call can take to be reconnected.

LPP route optimisation: navigation technologies so far have focused exclusively on physical route optimisations. As the role of connectivity gains more significance, navigation algorithms could use LPP forecasts for route optimisation purposes. This becomes increasingly relevant in an industrial context. Users adjust their behaviour over time, avoiding problem spots when they are known. Autonomous devices such as Automated Guided Vehicles, in an industrial context, do not have this awareness. With the help of LPP, such autonomous devices can become part of the closed-loop optimisation process.

FUTURE USE CASES FOR LPP

Whilst LPP is currently at the trials phase, valuable use cases are emerging. The remainder of this paper outlines future use cases that would benefit from real-time network performance forecasts like those generated by the AI engine at the core of LPP.

MOST NETWORK TRAFFIC IS TOLERANT TO AT LEAST SOME DELAY

Network traffic can be broadly categorised as:

- Delay tolerant, highly mobile
- Delay tolerant, stationary traffic
- Delay sensitive traffic

By predicting network performance in time and space, LPP is able to shift delay tolerant traffic in time or space in order to smooth out peak demand and improve overall network utilisation.

A significant proportion of mobile traffic demand is tolerant to delay. Video content, bulk software downloads, app and firmware updates for smartphones, IoT devices and connected vehicles are all applications which can tolerate some degree of time shifting. In the case of a video this might only mean that traffic can be brought forward, to increase the buffer, or delayed, to run it down, for only a few seconds. App updates or bulk software downloads could be delayed until the early hours of the morning or even a number of days.

Shifting mobile traffic demand in space is also a key feature of LPP. Pre-buffering a portion of video before a user enters an area of poor coverage or a busy town square is a possibility, for example.

Mobile video traffic accounted for 60% of mobile data in 2018 and Ericsson forecasts that this will rise to 74% in 2024[2]. If techniques such as LPP can be successfully applied to such

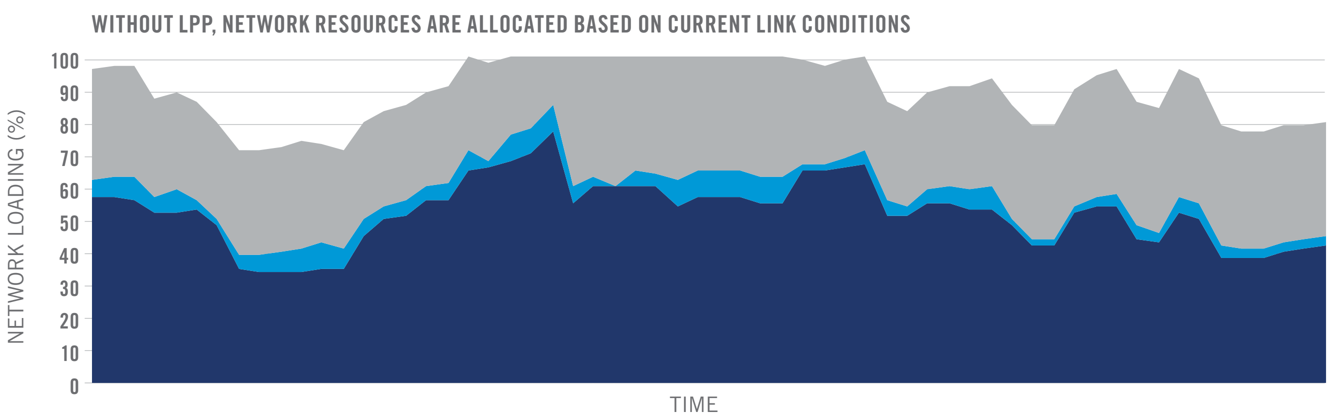

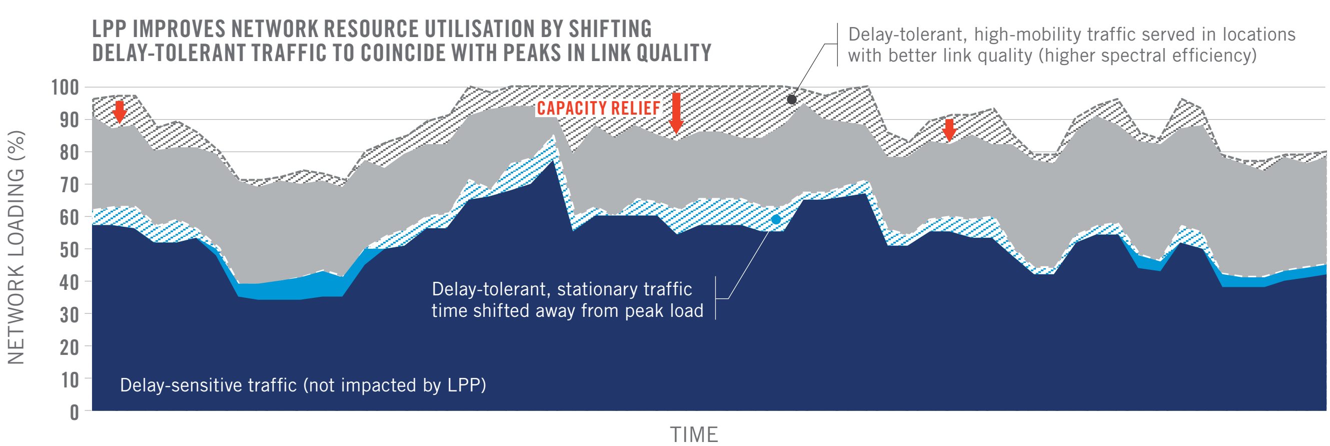

a large portion of network traffic, operators can use this to differentiate on Quality of Service or alternatively rationalise infrastructure deployment plans, in particular small cell roll out.

Whilst video content is the primary source of delay-tolerant traffic. A measurable impact could still be felt for other use cases, such as:

- Video gaming content in a hybrid-edge cloud architecture[3].
- Bulk software and high definition mapping downloads for connected and autonomous vehicles, drones and bots.
- File sharing services, such as Dropbox and Microsoft Sharepoint.

DELAY SENSITIVE TRAFFIC IS NOT DIRECTLY ADDRESSED BY LPP BUT MAY STILL BENEFIT

Today's 4G networks and IP networks typically work on a best efforts basis. When a particular high-reliability service is required, this is typically achieved by guaranteeing a tranche of network resources or providing priority access over other users without regard. The US emergency responder network, FirstNet, uses a combination of these approaches with a small guaranteed tranche of spectrum and priority over other users if demand overspills into the regular network[4].

5G requires a truly dynamic approach to network management to support a much broader range of use cases which demand ultra reliable and/or low latency communications. For example:

- Remote control of drones out of line of sight will require a constant connection
- Cloud gaming on mobile devices
- Live voice and video calls

Part of the benefit of the LPP approach is that in better managing the bulk of traffic on the network, an operator is able to more confidently guarantee the performance of a network slice dedicated to a high-reliability or low-latency use case.

The ability to forecast and communicate network conditions can also help users (or applications on their behalf) modify their behaviour to account for congested or otherwise suboptimal network conditions. For example, whilst a voice call or video call is not *time* delay tolerant, it could be tolerant to a change in location. For example, an autonomous vehicle could be slightly re-routed to maintain connection to a conference call, or the call could be held for a short period through a small section of poor coverage without dropping the connection.

The improved ability to guarantee the performance of critical services that this allows would further enhance the opportunity for operators to address use cases previously served almost exclusively by private networks. Examples of such opportunities include public safety, transport and organisations operating in hostile environments.

OPERATORS CAN MONETISE THE BENEFITS OF LPP TO SUBSCRIBERS IN A NUMBER OF WAYS

Experience shows that retail subscribers are unlikely to be willing to directly pay a premium for a higher quality of service, although they may be less likely to churn. In many cases churn reduction could be more significant to operator revenues than a straightforward ARPU premium.

The willingness of OTT providers to invest in their own CDN networks to guarantee their subscribers quality of experience suggests that there may be alternative mechanisms for monetising this differentiated service.

We expect part of the emerging 5G world is the provision of connectivity as part of a product or service. For example, Audi Connect makes use of an MVNO to provide infotainment services directly to its owners, rather than relying on their own mobile subscription. As eSIMs are more widely adopted in consumer goods such as wearables, particularly those that cannot rely on a domestic WiFi connection, manufacturers are increasingly likely to bundle connectivity for the benefit of consumers quality of experience, and their own data analytics. This is even more likely to be the case for connected B2B products and services such as the machinery and robots that will make up the Industrial Internet of Things.

So, whilst premium, industrial and enterprise connected devices will benefit from a high quality LPP enabled connectivity solution, corporate buyers are also in a position to make an informed purchasing decision in a way that retail consumers are not. Corporate buyers are able to invest in a thorough analysis of network performance and demand Service Level Guarantees

2 Ericsson Mobility Report, https://www.ericsson.com/assets/local/mobility-report/documents/2018/ericsson-mobility-report-november-2018.pdf
3 In this model content and processing are distributed between the user client, an edge server and a cloud server
4 FirstNet, https://firstnet.gov/sites/default/files/Top_Ten_FAQs_180514.pdf

in a way that retail subscribers are not. As this channel of selling connectivity services becomes more significant it will become important that operators can quantify their network performance and justify their pricing. This offers both an opportunity for premium pricing for operators that can provide a consistent high-quality experience and a threat to those who cannot demonstrate that their service lives up to its promise.

AS OPERATORS EXTEND THEIR ASSETS BEYOND CONNECTIVITY TO STORAGE AND COMPUTE, AUTOMATED EDGE MANAGEMENT ALSO BECOMES IMPORTANT

For mobile operators to live up to the promise of 5G connectivity, edge computing deployments will be critical. With potentially hundreds of thousands of "edges" (cell towers, campus, buildings, etc.) each with a constrained amount of storage and processing availability, managing the edge IT infrastructure will become as important and as challenging as managing the connectivity infrastructure itself.

By combining a forecast of both the service layer and network layer traffic loads, provided by LPP, into an automated management system, operators would be able to maximise the utilisation and quality of service delivered by the edge IT infrastructure.

LPP OFFERS COST SAVING POTENTIAL THROUGH ENHANCED NETWORK PLANNING AND MANAGEMENT

We have outlined above the benefits of LPP in terms of improved services or reduced requirements to deploy and operate infrastructure. The LPP approach can also be used as a more general predictive engine for network planning and maintenance.

- **Intelligent network management services:** using the predictions provided by LPP to autonomously manage transmitter power, antenna direction and tilt, and other services that might otherwise require direct NOC engineer intervention.

- **Network degradation control:** Identifying instances of poor performance and troubleshooting e.g. identifying when equipment should be repaired or replaced.

- **Efficient small cell backhaul planning:** Small cells are primarily designed to enhance coverage and/or capacity. For this reason, their locations are dispersed to cover discrete hotspots. This is challenging for both bandwidth and flexibility of the backhaul network. Using LPP as a predictive engine could help design a more efficient infrastructure deployment in terms of both small cells themselves but also the associated backhaul and power requirements.

- **Backhaul relief planning:** As described above, small cells have complex backhaul requirements, characterised by potentially peaky throughput. Operators can reduce operating costs by powering down links which are expected to require low usage at particular times before powering them back up at peak times.

CASE STUDY: SK TELECOM IS COLLABORATING WITH INTEL TO TRAIN LPP TO CONDUCT AUTOMATED NETWORK MANAGEMENT

As an early adopter of AI in telecommunications, SK Telecom has developed its own in-house AI-powered network operation system. The TANGO OSS (T Advanced Next Generation Operational Supporting System) uses big data analytics and machine learning capabilities to automatically detect network issues, and assist engineers in troubleshooting problems and optimising performance. TANGO OSS also provides for open-APIs which integrate inputs from services, such as Intel® LPP.

TANGO OSS was initially deployed on SK Telecom's network in late 2016, before extending across the mobile network in 2017.

SK Telecom and Intel's joint innovation programme has enabled the two companies to collaborate on leveraging LPP: a solution which builds on the lessons learned from TANGO OSS to develop a tool which draws on rich datasets to provide predictive abilities.

The partnership has enabled Intel's proof of concept algorithms to be trained using real network data, moving LPP closer to deployment readiness and providing the team with the insight required to understand how to adapt the technology to a specific customer's use case and network.

SK Telecom and Intel expect to move to large scale trials and initial deployment of LPP over the course of 2019 and will share progress in the coming months.

About Intel

Intel is known for its processors but is also a major investor in emerging machine learning technology. Link Performance Predictor (LPP) has emerged from this ongoing research at Intel's facilities. The development team are continuing to work to improve and optimise LPP based on a broad range of network conditions and use cases.

To discuss your potential use case further, please contact:

Jon Devlin, Business Development
jonathan.s.devlin@intel.com

Jonas Svennebring, Software Architect
jonas.svennebring@intel.com

Wooram Kim, Project Manager
alex.kim@intel.com

HeeSang Seo, Project Manager
daniel.seo@intel.com

About SK Telecom

SK Telecom is the largest mobile operator in Korea and known for its significant role in the advancement of mobile telecommunication technologies. SK Telecom is collaborating with Intel on LPP, including trialling the technology for intelligent network management services and network management control. This will support SK Telecom's ambition to provide the connectivity services required to enable the Fourth Industrial Revolution.

For more details, please contact:

Dong Kee Lee, Team Leader
dongkee.lee@sk.com

Woohyun Nam, Project Manager
namwh@sk.com

Keun Hyun Kim, Project Manager
keunhyun.kim@sk.com

Daniel H. Chae, Project Manager
dani75.chae@sk.com

About Cambridge Consultants

Cambridge Consultants is a world-class supplier of innovative product development engineering and technology consulting. We work with companies globally to help them manage the business impact of the changing technology landscape.

With a team of more than 800 staff in the UK, the USA, Singapore and Japan, we have all the in-house skills needed to help you – from creating innovative concepts right the way through to taking your product into manufacturing. Most of our projects deliver prototype hardware or software and trials production batches. Equally, our technology consultants can help you to maximise your product portfolio and technology roadmap.

We're not content just to create 'me-too' products that make incremental change; we specialise in helping companies achieve the seemingly impossible. We work with some of the world's largest blue-chip companies as well as with some of the smallest, innovative start-ups that want to change the status quo fast.

Cambridge Consultants is part of the Altran Group, a global leader in innovation. www.Altran.com

For more information, please contact:

Michal Gabrielczyk, Senior Consultant – Technology Strategy
michal.gabrielczyk@cambridgeconsultants.com

Sandro Grech, Senior Consultant – Technology Strategy
sandro.grech@cambridgeconsultants.com