Ericsson White Paper
GFTL-20:001402
February 2022

ERICSSON

6G – Connecting a cyber-physical world

A research outlook toward 2030

Content

Introduction

5G expansion is continuing throughout the world, with networks providing new communication capabilities and services that are set to transform society. The next wave of development is now taking place through 5G Advanced, with improved capabilities in the areas of enhanced Mobile Broadband (eMBB), ultra-reliable low latency communication (URLLC), and massive Machine Type Communication (mMTC).

There is no doubt that the ongoing societal transformation will give rise to challenges that 5G will be unable to meet. In 2030, society will have been shaped by 5G for 10 years, with lessons having been learned from 5G deployment, and new needs and services appearing. Even with the built-in flexibility of 5G, we will see a need for expanding into new capabilities [1]. This calls for further evolution—following the pull from society's needs and the push from more advanced technological tools becoming available—that must be addressed for the 6G era when it comes.

Future networks will be a fundamental component for the functioning of virtually all parts of life, society, and industries, fulfilling the communication needs of humans as well as intelligent machines. To make the best out of this situation, both the industry and research community should work together toward a common vision.

Four main drivers with corresponding challenges are emerging for the 6G era: trustworthiness of the systems at the heart of society, sustainability through the efficiency of mobile technology, accelerated automatization and digitalization to simplify and improve people's lives, and limitless connectivity meeting the demands for intensifying communication anywhere, anytime, and for anything.

To meet these future challenges, 6G needs to continue to push beyond the technical limits of 5G, moving toward critical services, immersive communication, and omnipresent IoT. In addition, entirely new capability dimensions should be explored integrating compute services and offering functionality beyond communication such as spatial and timing data.

In this white paper, a vision of the 6G-powered world of 2030 is outlined, with the research focusing on what future networks should be able to deliver and what candidate technologies should be developed to get there.

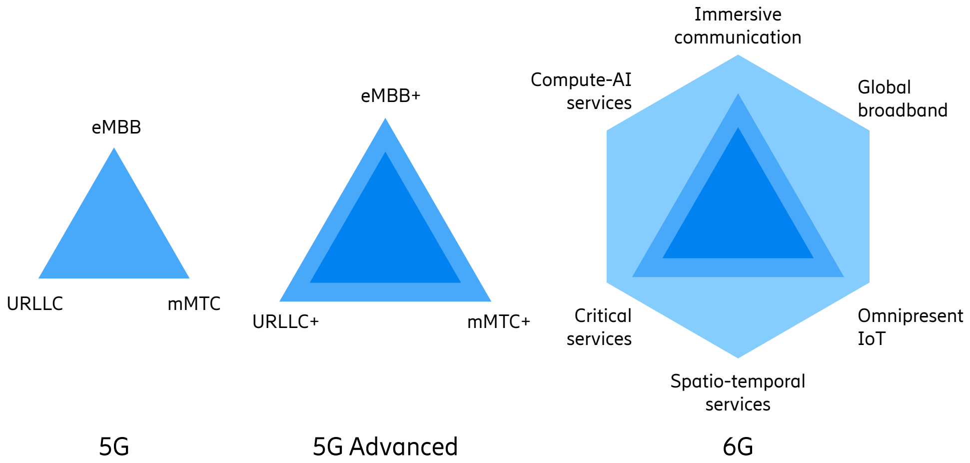
eMBB
URLLC
mMTC
eMBB+
URLLC+
mMTC+
Immersive communication
Compute-AI services
Global broadband
Critical services
Omnipresent IoT
Spatio-temporal services
5G
5G Advanced
6G

The cyber-physical world of 2030

The society of 2030 is expected to have transformed around increasingly advanced technologies, where networks act as the communication and information backbone, allowing communication to take place anywhere and at any time.

As wireless connectivity becomes an integrated, fundamental part of society, trust in the data delivered through connectivity as well as in the connectivity itself, along with data services and the compute platform functionality, will become even more important. Society should be able to rely completely on networks to deliver critical services and to ensure the integrity of the delivered information. People, as well as industries, must be able to rely on verified identities while enjoying full privacy.

Sustainability is of utmost importance, and all sectors of society need to work toward the United Nations Sustainable Development Goals (SDGs) [2]. Wireless networks already play an important role in achieving these goals, and there is clear potential to further accelerate their contribution in enabling better efficiency in the use of resources and supporting new ways of living, making them a tool for positive change.

With artificial intelligence (AI), it is possible to optimize and simplify many processes and improve operations by reducing the need for human participation and supervision. As a result, a dramatic increase in the use of AI to further optimize efficiency in society and simplify people's lives should be expected. To enable this, networks need to be based on a data-driven architecture using massive data to support AI across systems and be designed for the highest levels of security and explainability.

Today, there is a large increase in highly demanding applications for which very low latencies and very high data rates are required to enable their functioning such as virtual, augmented, and mixed reality as well as remote control of sensitive operations. As 2030 approaches, this trend can be expected to continue with even higher demands being placed on the performance that networks should deliver.

Trustworthiness	Trusted communication and computing for industry and society relying on critical information
Sustainable world	Communication and network as part of and enabler for sustainable development
Simplified life	Massive use of AI across systems for optimal assistance and efficiency
Application demands	Extended and new services requiring extreme connectivity performance

6G paradigm shifts

Addressing the challenges of the future also implies that there must be a shift in some basic network paradigms. The shifts that must take place are:

From secure communication to trustworthy platform	From energy efficiency to sustainable transformation	From manually controlled to learning networks	From physical and digital worlds to a cyber-physical continuum
From data management to data ownership	From terrestrial 2D to global 3D connectivity	From predefined services to flexible user centricity	From data links to services beyond communication

- from secure communication to trustworthy platforms—expanding the scope from protecting data to ensuring the end-to-end service delivery in the relevant scenarios
- from data management to data ownership—ensuring control and privacy of personal and critical digital assets toward third parties
- from energy efficiency to sustainable transformation—resource-efficient networks that impact society and enable environmental footprint reduction through effective digitalization

- from terrestrial 2D to global 3D connectivity—aiming at full digital inclusion, reaching for limitless connectivity everywhere including rural land, sea, and even air areas
- from manually controlled to learning networks—using intelligence and data throughout networks to shift the focus from instructing the system how to achieve goals, to providing the system with goals to achieve
- from predefined services to flexible user-centricity—instead of predefining services and interfaces, a flexible network that should adapt to the needs of users and allow applications to influence
- from physical and digital worlds to a cyber-physical continuum—the network platform should not only connect humans and machines but be able to fully merge realities to allow seamless interaction and immersive experiences
- from data links to services beyond communication—expanding the role of networks to deliver services for a broad range of purposes as a versatile information platform

Transforming society through technology

6G makes it possible to move in a cyber-physical continuum, between the connected physical world of senses, actions, and experiences and its programmable digital representation. The network provides intelligence, limitless connectivity, and full synchronization of the physical and digital worlds. Vast amounts of sensors embedded in the physical world send data to update the digital representation in real time. Actuators in the real world carry out commands from intelligent agents in the digital world. It becomes possible to trace back and analyze events, observe, and act in real-time, as well as to simulate, predict, and program future actions. Compared to the metaverse—a VR/AR world where avatars interact—the cyber-physical continuum provides a close link to reality, where digital objects are projected onto physical objects that are represented digitally, allowing them to seamlessly coexist as merged reality and enhance the real world.

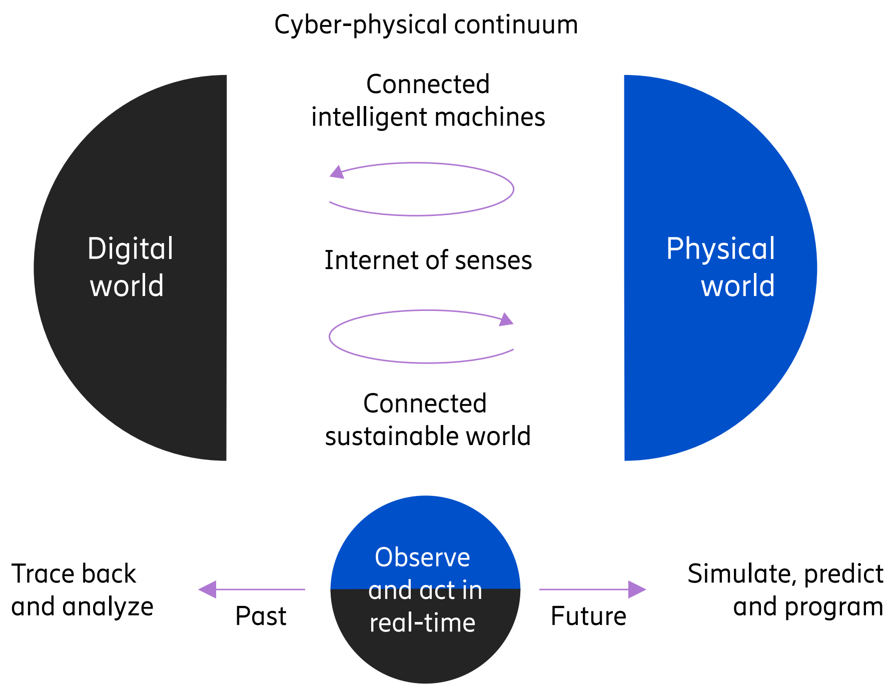

Triggered by the four drivers described above, new application areas will appear, calling for new capabilities in the networks of the future. It is equally important to address the future of existing applications and close the remaining digital gaps.

A digitalized and programmable world can deliver interactive 4D maps of whole cities that are precise in position and time and can be simultaneously accessed and modified by large numbers of humans and intelligent machines for detailed planning of activities. Such cyber-physical service platforms can issue commands to large-scale steerable systems, like public transport, waste handling, or water and heating management systems, achieving higher levels of resource efficiency, better control, and increased resilience.

The advent of precision healthcare, enabled by miniature nodes measuring bodily functions and devices issuing medications and physical assistance, will be supported by a continuously analyzed digital representation online. Such a high integration of technology in people's lives emphasizes the importance of trustworthiness through availability, security, and data privacy. It also requires new types of devices that can be safely embedded virtually anywhere and that are maintenance-free, using efficient and distributed processing and management, and communicating securely in local body networks.

Real-time 4D maps are also needed to manage the intense traffic of future cities with autonomous vehicles on the ground and in the air. A network sensor fabric, where accurate measurements and world data are aggregated from sensing base stations and on-board vehicle sensors, and then shared together with trajectories, can be used to guide safe, clean, and efficient transport.

An automated society would harvest the benefits of AI assistance for improving people's welfare and simplifying their lives. For instance, collaborative AI partners could perform many challenging tasks involving manual labor more safely and efficiently, assisting in industries as well as in our homes, acting autonomously, and adapting to human action. Such high-trust cyber-physical systems can smoothly interact with groups of humans and other intelligent machines, requiring extreme reliability and resilience, precise positioning and sensing, low-latency communication, and AI trust and integration. On the personal level, intelligent identity and preference handling will assist people in everyday life, managing interactions with and adapting the connected world around them in line with their preferences.

Building a sustainable world requires huge efforts throughout society, with networks ensuring digital inclusion on a global scale. This includes diverse elements, such as the support of smart automation services everywhere on the planet, connectivity for global sensors monitoring the statuses of forests and oceans, resource-efficient connected agriculture, access to digital personal healthcare for everyone, and access to high-end services for institutions such as schools and hospitals everywhere. Through the global, end-to-end life-cycle tracking of goods, autonomous supply chains can accelerate a full circular resource economy. Digital-asset tracking can reduce waste and automatize recycling. Taken together, this requires truly global coverage with excellent energy-, material-, and cost-efficiency, embedded autonomous devices and sensors, and a network platform with high availability and security.

Immersive communication will deliver the full telepresence experience, removing distance as a barrier to interaction. Extended reality (XR) technology [3] with human-grade sensory feedback requires high data rates and capacity, spatial mapping from precise positioning and sensing, and low latency end-to-end with edge cloud processing. One example will be the ubiquitous use of mixed reality in public transport, offering separate virtual experiences for each passenger, enabling them to run virtual errands, get XR guidance, and have games overlaid on the physical world. Going further still, communication will approach a fully merged reality where senses and holograms can be translated across physical and digital worlds. Personal immersive devices capable of precise body interaction will allow access to experiences and actions far away, ensuring an immersive perception, to support people's communication needs even better—the importance of which has become especially clear during the COVID-19 pandemic—and, at the same time, add completely new communication modes, with strict control over access and identities.

6G: the future network platform

Increasing expectations set a clear target for the industry and research community—6G should contribute to an efficient, human-friendly, sustainable society through ever-present intelligent communication.

Needed capabilities

To serve as the platform for a vast range of new and evolving services, the capabilities of future wireless access networks need to be enhanced and extended in various dimensions compared to the networks of today. This includes classic capabilities, such as achievable data rates, latency, and system capacity, but also new capabilities, some of which may be more qualitative in nature. It should be noted that the capabilities of the future wireless networks should not only match currently envisioned use cases but should also enable future services that have not been envisioned yet.

Starting with the classic capabilities, future networks should enable higher **achievable** data rates and lower latency in **all relevant scenarios**. This includes the possibility to provide several hundred gigabits per second and end-to-end sub-millisecond latency in specific scenarios. Equally or perhaps even more important is the possibility to provide high-speed connectivity with predictably low latency and a low jitter rate.

The future wireless-access networks should be able to serve an exponentially growing traffic demand in a cost-efficient way. Higher spectral efficiency of basic radio access technology is one component of this, with access to additional spectrum naturally being another. Even more important, though, is enabling the possibility for the truly cost-efficient deployment of very dense networks.

There is a need to continue the expansion of wireless communication toward full global coverage—closing the digital divide for remote areas—while supporting a dramatically higher number of devices that will be embedded throughout society. A critical component of this is to ensure that the overall cost for both users and service providers is at a sustainable level.

High network energy performance played was an important requirement in the development of 5G, and it will be even more important for future wireless access solutions [4]. It is critical that the expected massive increase in traffic will not lead to a corresponding increase in energy usage. An acceleration in traffic should not mean accelerated energy usage. Also, the energy usage should be close to zero when there is no traffic within a node.

As wireless networks increasingly become critical components of society, resilience and security capabilities are crucial. The networks must be able to provide service when part of the infrastructure is disabled due to natural disasters, local disturbances, or societal breakdowns, and they must offer robust resistance against deliberate malicious attacks.

In terms of trustworthiness, the networks should be able to leverage new confidential computing technologies, improve service availability, and provide enhanced security identities and protocols with end-to-end assurance.

These networks will need the capabilities of dependable compute and AI integration, infrastructure enabling distributed applications and network functions to be swiftly developed and deployed, and services for data and compute acceleration, which can be delivered throughout the network with performance guarantees.

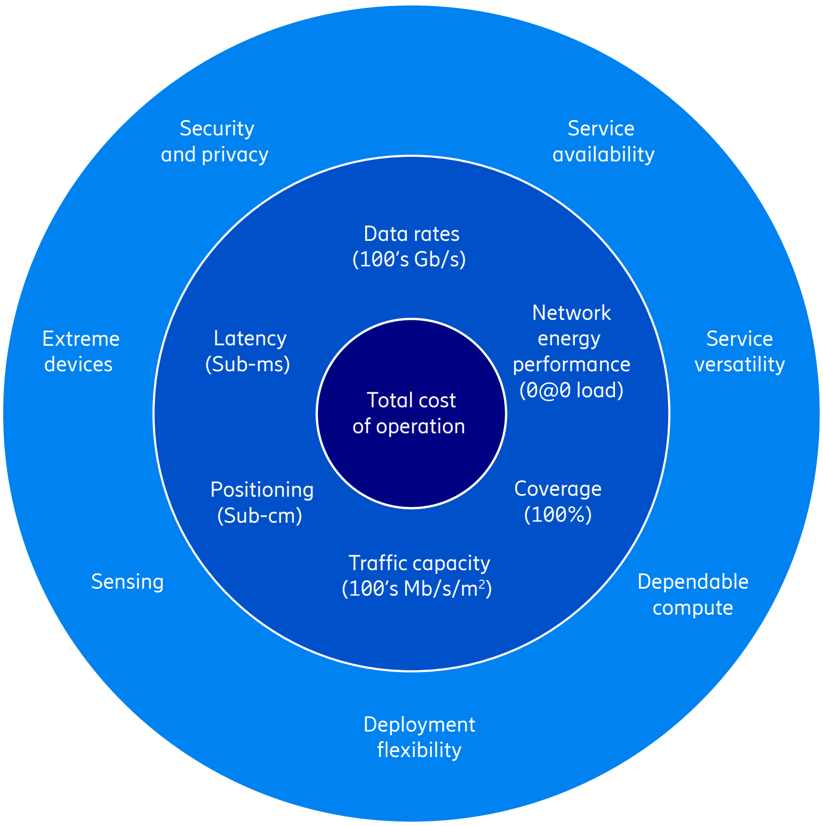

Finally, to power the full digitalization and automation of society, networks need high-precision positioning and detailed sensing capabilities from their surroundings. Sensing is a new type of capability, based on analysis of how radio wave propagation is affected by the environment. For example, microwave links are affected by rainfall—information that is valuable for weather forecasting. Radio signals can also be actively transmitted for sensing, allowing a radar-like function to be provided by the communication network. Reusing cellular systems for sensing can result in both a more cost-efficient sensing system and broader coverage than what can be provided by dedicated sensing systems. Sensing can be used, for instance, to model the environment, detect road traffic, or set off alarms if a person enters a restricted area in a factory hall. Future networks need to use radio resources efficiently for both communication and sensing. Scalable mechanisms for distributing the results, AI-based interpretation of results, and security mechanisms to ensure the privacy of the information are also needed.

Foundations of 6G

For the 6G networks of 2030, a broad range of promising technologies should be considered. The study of these potential elements of 6G will be a key topic of research in the coming years.

Taken together, the elements of 6G will form one seamless system, having all the needed capabilities to empower the vision of ever-present intelligent communication connecting a cyber-physical world. With a foundation of trustworthy systems and a highly efficient compute fabric with built-in cognition capacities, the networks of the future will deliver limitless connectivity for upcoming applications and services. This will make 6G a broad platform for innovation and the information backbone of society.

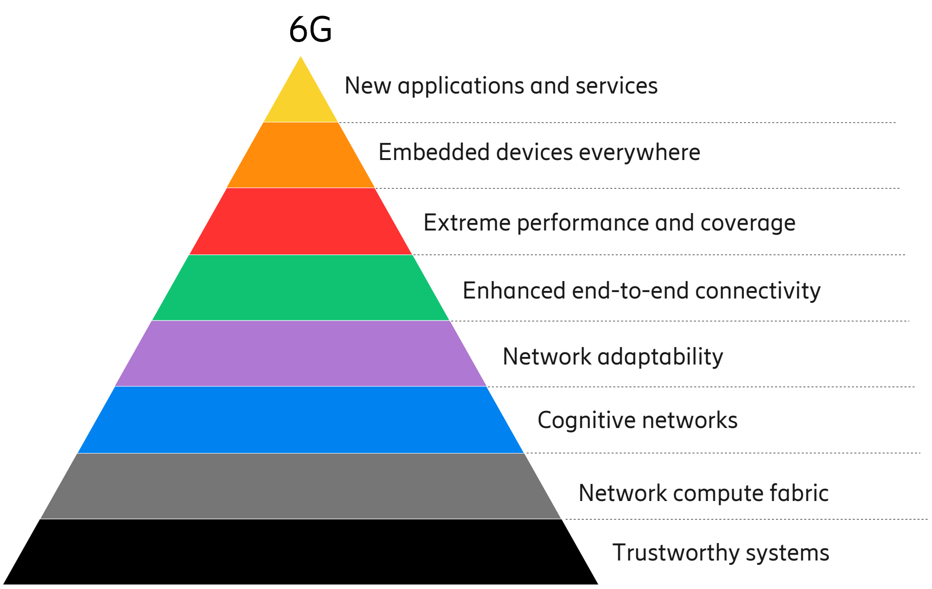

Technology elements of 6G

Network adaptability

By increasing the adaptability of networks, several key efficiencies can be addressed. These might be related to the cost of deployments, energy consumption, network development and expansion, or management and operations.

Dynamic network deployment

Mechanisms to ensure that dynamic network deployment will be key to supporting the cost-effective deployment of high capacity, resilient networks in the future. This will make the service provider more agile when handling new business opportunities and new emerging use cases. The key challenge is to seamlessly integrate traditional service provider-deployed network nodes with complementing ad-hoc, temporary, user-deployed, mobile, or non-terrestrial nodes.

The possibility for multi-hop communication—already partly introduced in 5G through integrated access backhauling (IAB)—will be an important component to enabling such dynamic network deployments. It is expected that this will further evolve, ensuring seamless multi-hop wireless connectivity with low costs and high flexibility. This will also partly erase the distinction between wireless access links to devices and wireless backhaul links between network nodes, creating a unified framework for wireless connectivity.

A factor that is common to all future deployment scenarios is the requirement for a superior transport network to be flexible, scalable, and reliable, in order to support demanding 6G use cases and novel deployment options, such as a mixture of distributed radio access networks (RANs) and centralized/cloud RAN. This is achieved by AI-powered programmability enabled by software definition, multi-service abstraction/virtualization on heterogeneous networks, and closed-loop automation to keep transport networks flexible and manageable.

Device and network programmability

Previous generations of cellular networks have relied on clearly specified device behaviors controlled by network configurations. The key limitation here is that new features cannot be applied to legacy devices, limiting the speed of development.

Device behaviors can be made more programmable, making devices more future-proof and ready to support more advanced network functionalities by replacing hardcoded device behaviors with a more programmable environment (for example, defining them by different application programming interfaces—APIs). This, in turn, would enable networks to be more programmable, since it would now be possible to fundamentally change both the networks and the devices, enabling new functionalities (for example, allowing service providers to download AI models to both the devices and the networks, optimizing the overall network performance or customizing the device behavior targeting specific vertical use cases). Another aspect is that this could lead to faster feature development and time to market, faster bug fixing, and more DevOps-type operations.

Network simplification and cross-RAN/CN optimizations

With the expectation of networks becoming a more integral part of society comes the requirements of higher availability and resilience. Over the years, however, networks have continuously grown in functional richness as well as complexity. This has led to multiple network components supporting many different functions and sometimes addressing similar (or identical) problems. Future deployments will be less node-centric and both RAN and core networks (CNs) will have more common platforms. This removes some of the reasons to duplicate functionalities, such as having RAN rely on the CN as a data store for idle devices. Consequently, it is important to revisit some architecture assumptions behind today's functional separation between RAN and CN.

A smart choice when it comes to the right set of RAN and CN functions and interfaces is needed to provide the best performance, use cases, and deployment versatility while at the same time keeping development efforts and network operations manageable. A set of multi-vendor interfaces needs to be selected carefully to ensure openness in networks and the ecosystem while minimizing system complexity, ensuring development agility and a robust and resilient network.

Enhanced, end-to-end connectivity

Future applications need to leverage high-performance connectivity, fulfilling required bandwidth, dynamic behaviors, resilience, and further demands. Network capabilities need to be available end-to-end and match the evolution of applications and internet technology. This will affect, for instance, application–network collaboration, resilience mechanisms, the evolution of the end-to-end transport protocols, and ways to deal with latency.

Network collaboration

Applications and networks can benefit from collaboration to ensure that the most suitable networking services are provided. The increased need for protected communications implies that any collaboration needs to be explicitly agreed upon, with both parties benefitting from and consenting to it.

Resilience

Network resilience will need to be addressed from different perspectives. Applications that demand resilience, both for their connectivity and their end-to-end communication, need to be supported. Similarly, the necessary internet infrastructure needs to be available, resilient, and resistant to commercial surveillance. A distributed architecture ensures that not all information (and not all risk) is centralized among a few parties.

Evolved protocols

The recent rapid evolution of web and transport protocols has resulted in the internet protocol stack becoming easier to change (for example, it is now possible to update transport protocols without impacting operating system kernels). At the same time, it is expected that future communications will employ more multi-access technology and applications to come with even stricter requirements. This is an opportunity to build solutions that can handle multi-path communications, resilience, and congestion control in mobile networks more efficiently.

Predictable latency

Experience has shown that many of the (initial) use cases with stricter latency requirements often have a maximum latency that they tolerate. Achieving predictable latency will open opportunities for testing additional use cases and support both distributed and more centralized deployment models.

Extreme performance and coverage

The future wireless access solution must provide truly extreme performance in a multitude of capability dimensions and all relevant scenarios in order to enable future in-demand services at acceptable costs. This includes, for example, extreme data rates and latency performance when so required, extreme system capacity to be able to deliver the services to a high number of users, and truly global coverage of the wireless access. The key to enabling dense deployments with extreme system capacity in a cost-effective way is to introduce packet fronthaul and new wireless transport technologies, such as relay and mesh networking, free-space optics, and further integrated access and backhaul.

Spectrum

Spectrum is—and will continue to be—an obviously essential resource for wireless connectivity. Access to additional wideband spectrum as well as efficient utilization of the existing spectrum is of critical importance, and both licensed and unlicensed spectrum are of interest.

The lower frequency bands (up to about 6 GHz) are currently used by 4G/5G and will remain important in the 6G era, especially to provide wide-area coverage for 6G services. Since very little new sub-6 GHz spectrum is expected to be made available, it is essential that a 6G radio access technology will be able to share lower-frequency spectrum with previous generations. The millimeter wave frequency bands in the 24 GHz to 52 GHz range, pioneered by 5G and likely to soon be extended up to 100 GHz, will naturally be used by 6G as well.

The 7–24 GHz range is currently being used for other purposes than cellular communication but can be exploited for 6G by deploying advanced sharing mechanisms. Above 100 GHz, there are opportunities for relatively large amounts of spectrum, but, given the very challenging propagation conditions, it is mainly of interest for very specific scenarios requiring extreme traffic capacity and/or data rates in a dense network deployment condition.

Non-terrestrial access

Extending the conventional terrestrial access to also include non-terrestrial (NT) access components is a tool to realize truly global coverage for future wireless connectivity. Such complementary NT access components may be provided by different means, including, for example, drones, high-altitude platforms (HAPS), and/or low-Earth orbit (LEO) satellites. These mobile NT nodes should be integrated parts of the overall wireless access solution as an extension of the terrestrial network, providing seamless coverage truly everywhere.

Multi-connectivity and distributed MIMO

In order to enhance robustness and performance as well as ensure more consistent quality in wireless connectivity, multi-point connectivity is expected to become common in the future. Already today, technologies such as multi-radio, dual-connectivity, and multi–point transmissions are available for 5G, but it is expected that they will expand further. This expansion might include, for example, massive multi-connectivity on the physical layer, where devices have simultaneous physical links to a large number of tightly coordinated network transmission points (known as distributed MIMO). Another possibility is multi-RAT connectivity, where devices have simultaneous connectivity to a network using different radio access technologies to improve robustness or to provide different simultaneous services in a more optimized way.

Embedded devices everywhere

Future services will require connectivity everywhere and in everything. 6G networks can support trillions of embeddable devices and provide trustworthy connections that are available all the time.

Zero-energy devices

Today's massive machine-type communication provides data rates of up to a few hundred kilobits per second, serving applications such as remote meter reading. Although their battery life can be up to 10 years in some cases, battery replacement or charging limits the applicability of these devices. Energy harvesting—where a device's energy is obtained from ambient energy in the form of light, vibrations, temperature differences, or even radio waves—provides the possibility for devices to not need a battery replacement or charging. The amount of energy that it is possible to harvest is typically very small, however, implying that extremely energy-efficient communication protocols need to be developed. Given the minuscule amounts of energy available, the amount of information that can be transmitted will be small—in many cases, only a couple of bytes per hour. For applications such as asset tracking, however, this is sufficient, and radio-based technologies could be a more appealing choice than the current solutions, such as the optical reading of barcodes, and would facilitate communication with items out of direct sight.

Immersive interaction devices

In future, users will be able to have a more immersive experience where they can interact naturally with the digital world, with the help of on-body devices, such as smart gloves, skin sensors, and so on. Users will interface with virtual objects, which often requires accurate positioning - as in the case of dropping virtual ice cubes into a virtual drink - and experience the virtual object's updates in real time with all the user's senses, requiring sub-millisecond latencies. Brain-computer interface (BCI) devices could further enhance the experience by capturing and securely sharing users' intentions to adapt to the network rendered virtual objects. Networks will also aid in synchronization between such objects and sensory stimuli beyond the visual such as audio, touch, and so on. Furthermore, trustworthiness aspects such as the verification of user IDs to protect vulnerable users from inappropriate content and contact would need to be addressed.

Cognitive networks

To realize future networks and operate a large number of versatile services without accelerating cost and complexity, the level of network intelligence must be raised. The resulting cognitive networks will help improve energy efficiency, optimize performance, and ensure service availability. It is expected that this will occur in two ways: in optimizations that are difficult to achieve with traditional algorithms, where AI machine learning (ML) can support, and in evolving the operations systems to handle most of today's system management tasks autonomously, where AI machine reasoning (MR) can play a vital role.

Intent-based management

Humans will be able to control what systems do by stating operational goals in the form of intents. This intent-based, automated management requires a higher level of abstraction in the human-machine interface as well as the systems' ability to interpret and reason around such goals. There is a need to understand abstract knowledge and draw conclusions from existing knowledge and data sets using MR techniques. Knowledge and experience will be gathered, both from humans and analytics algorithms, and stored in a common knowledge base. These varying elements would then be used by a cognitive network to understand different situations, identify suitable corrective measures, and plan the best course of action for their implementation in the network.

Autonomous systems

Such an approach also implies that the system becomes more and more autonomous. A cognitive system requires native capabilities to adjust to its environment, constantly observing and learning from previous actions. Lessons from operations and service performance are fed back in short cycles or in near real time to improve configurations, processes, and software. Within the network logic, a continuous improvement in algorithms will be seen driving runtime decisions distributed across physical locations and logical functions. This continuous optimization will make the system much more dynamic compared to today's system. Intelligence, in different forms, will be available all over a geographically distributed network.

Explainable and trustworthy AI

An autonomous system can only be successful if it is trusted by humans. This involves several aspects. Firstly, the system needs to be able to explain its actions and why it ended up in its current state. Secondly, the intelligent system should be technically robust, even under various disturbances and attacks; consider its social environment; act ethically, respecting the right principles and values; and act in accordance with all applicable laws and regulations [5]. Thirdly, the system must involve humans when needed.

Data-driven architecture

Intelligence involves making decisions based on facts or data, and with more data available, better decisions can be made. Data-driven architecture is the infrastructure for AI algorithms that makes decisions. Such infrastructure supports data pipelines that take care of moving, storing, processing, visualizing, and exposing data from inside service provider networks as well as external data sources in a format adapted for the consumer of the pipeline.

Network compute fabric

6G will bring all physical things into the realm of compute. It will act not only as a connector but also as a controller of physical systems—ranging from simple terminals to complex and performance-sensitive robot control systems, and augmented reality applications—offering computing intertwined with communication in a network compute fabric for high efficiency and dependability [6].

Service providers can utilize their assets by integrating compute and storage into increasingly virtualized networks to provide applications with maximum performance, reliability, low jitter, and millisecond latencies. The network-compute fabric will thus provide tools and services beyond connectivity, offering a pervasive, globally interconnected, flexible, and robust platform to diverse customer segments and verticals, featuring application hosting, seamless task portability, and compute abstractions.

Ecosystem enablers

Such a system can only be realized with the collaboration of a broad set of actors working in the same globally federated ecosystem. Network and cloud providers, application developers, service providers, and device and equipment vendors all have a role to play. Much of the interaction between the players will happen in software, where broker-less marketplace technologies will help the ecosystem to scale, featuring automated contract negotiation and fulfillment supporting sales, delivery, and charging operations. Ecosystem partnerships will also involve technical challenges of integrating services from the different actors. Efficient partnering can be facilitated by standards that ensure interoperability or by technologies that automate the handling of partner relationships. The right level of harmonization in the ecosystem is key to supporting scalability as well as innovation.

Dependable compute

Emerging use cases will require a combination of stringent real-time characteristics, such as low latency, high throughput, high reliability, and scalability. To meet these end-to-end performance requirements, 6G platforms will complement deterministic and reliable connectivity offerings with corresponding compute capabilities. Through a network compute fabric, the network will offer unified interfaces for the simplified deployment of distributed applications onto an integrated compute stack with dependable real-time properties for critical application tasks. For instance, developers will have access to in-network compute services realized through energy-efficient hardware acceleration technologies as well as the operating system and platform components optimized for real-time operations.

Unified, fluid computing

Applications developed to interact with physical reality need increased deployment flexibility. They benefit from highly distributed designs in order to be close to data sources and data consumers, such as sensors and actuators, for instance, in closed-loop control of mission-critical processes, and intelligent aggregation of large amounts of data. Furthermore, smartly splitting the processing between devices, network compute, and central clouds enables richer applications on extremely lightweight devices with limited power supply. This poses several new challenges to computing. New ways of combining, placing, and executing software are needed to meet application requirements even in the face of user mobility or failures.

For instance, dynamic computational offload, that is, moving applications tasks that are split off from device applications into network embedded compute, will benefit from a unified execution environment based on lightweight and portable runtime technologies, hardware-enforced isolation for secure task execution, and developer-friendly exposure and interfaces. As a result, applications can be deployed seamlessly across this fabric, spanning central cloud, network edge, and out to the devices. This makes 6G a true innovation platform.

Trustworthy systems

The ability to withstand, detect, respond to, and recover from attacks and unintentional disturbances is a cornerstone in designing trustworthy systems. The four important building blocks for trustworthy systems are the use of confidential computing solutions, secure identities and protocols, service availability, and security assurance and defense.

AI is expected to have a major impact on future technology evolution as well as security and to help in all of these four areas. At the same time, the trustworthiness of AI components is also important.

Confidential computing

Today's systems offer strong protection for data in transit, but data being processed or stored is less protected. To protect data being processed or stored, confidential computing is becoming a strong paradigm. In cloud computing, it provides hardware-based isolation for the processing of payloads, which a cloud provider cannot tamper with. It also allows remote cloud users to verify the isolated environments in which they want to place their payloads, and the verification and attestation procedures are carried out by the compute hardware itself, preventing a bypass by the cloud provider owning the hardware. The basis of these confidential computing features is part of the root-of-trust (RoT) mechanism.

Confidential computing also has the potential to enhance the security of network slices. Network slices can be cryptographically isolated from each other by combining data in transit protection mechanisms and confidential computing technologies to protect data being processed or stored. The path to secure identities and protocols depends on establishing trusted identities for infrastructure, connectivity, devices, edge, and network slicing functions. This can be enabled by means of RoT mechanisms for identities that are established for every physical component, a software function, and interface. The end goal is to create a system that offers privacy for all deployed software as well as the protection of data from unauthorized access.

Service availability

Service availability can be offered by paying attention to details that improve the reliability and resilience of networks as a whole. The radio link is a critical part of meeting availability requirements. Radio resilience can be improved by the provisioning of adequate capacity, redundancy of coverage, and the use of the diversity of connectivity and medium access control. Resource provisioning for critical services across RAN, transport, and core can be designed to allow variable grades of service and service guarantees in terms of meeting near-real-time deadlines in industrial scenarios or other critical control functions.

Another aspect of availability is the building of automated recovery mechanisms by analyzing and aggregating data from all parts of the communications system. This means that a distributed and hierarchical approach for improved observability of performance must be designed, with intermediate analytics, which can validate that requirements are being met on a real-time basis. Furthermore, AI has a role to play in integrating data-driven observability to infer the end-to-end validation of service availability. Real-time analytics based on AI can similarly offer the ability to improve network resilience to dynamic changes in traffic load and radio environments, thereby providing further assurance of the reliability of performance in relation to the needs of various network slices.

Security assurance

Today, security assurance and certification are receiving a lot of attention. For example, the EU Cybersecurity Act, adopted in 2019, establishes an EU framework for cybersecurity certification, boosting the cybersecurity of digital products and services in Europe. Current state-of-the-art security assurance schemes (for example, GSMA NESAS) are good tools for providing security assurance for a specific version of a product; however, some areas need further development. Enhancements for virtualization and cloud computing, continuous integration and continuous delivery processes, and AI need to be taken into account. An essential aspect to consider here is that security has a much wider interpretation than just product security, which is what security assurance schemes concentrate on today. In the future, they should be amended to increasingly consider all aspects of the system, including networks in operation. When creating security assurance schemes, it is important to establish well-defined requirements and processes that are accepted by all stakeholders. This is preferably achieved in line with global standards.

Conclusion

There is a strong upcoming need for communication and beyond technology on the 2030 horizon, with the transformations having been set in motion by 5G, and increasing expectations in society, accelerated by advancements in enabling technology, and moving toward new services and use cases that will improve people's lives.

Development is ramping up in formulating capability targets for the 6G era and investigating a range of promising technology components that may become part of a 2030 network platform. The key elements for this transformation will be the extreme performance of radio access, with network adaptability, and global as well as pervasive reach. Going beyond connectivity, 6G should become a trusted platform for intelligence, compute, and spatial data, encouraging innovation and serving as the information backbone of society.

This is the right time for advanced research on 6G technology into a network platform aimed at expanding the capabilities for the needs of 2030. Research collaborations like the Hexa-X project [7] and the Next G Alliance [8] are already advancing in technology and system design to enable the cyber-physical world and limitless connectivity. The journey toward future networks should naturally build on the strengths of 5G, which continue to evolve in 5G Advanced, and should be taken in collaboration with the academic community and other industry partners aiming at a globally aligned way forward, building between regional initiatives.

References

1. Five network trends towards the 6G era,
   https://www.ericsson.com/en/reports-and-papers/ericsson-technology-review/articles/technology-trends-2021
2. UN Sustainability Development Goals,
   www.undp.org/content/undp/en/home/sustainable-development-goals.html
3. The Internet of Senses report,
   www.ericsson.com/en/reports-and-papers/consumerlab/reports/10-hot-consumer-trends-2030
4. Exponential roadmap, exponentialroadmap.org/
5. Trustworthy AI,
   ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai
6. The network compute fabric, Ericsson Technology Review,
   https://www.ericsson.com/en/reports-and-papers/ericsson-technology-review/articles/network-compute-fabric
7. The Hexa-X project, hexa-x.eu
8. The Next G Alliance, nextgalliance.org

Authors

Gustav Wikström is a Research Leader at Ericsson Research Networks, driving toward the next generation of networks. He joined Ericsson in 2011 after completing his postdoctoral studies in physics and has worked with standardization, concept development, and performance evaluations for WLAN, 4G, and 5G. He is currently focusing on 6G vision, use cases, and services.

Patrik Persson is holding the position as 6G program manager director at Ericsson Research, having joined in 2007. Before focusing on driving the 6G vision and concept research activities Patrik has been involved in different areas including concept development for advanced antenna systems, 3GPP RAN standardization (4G and 5G), and proprietary evolution of LTE. Patrik holds a PhD (2002) and a docent degree (2011) in electrical engineering from the Royal Institute of Technology (KTH), Stockholm, Sweden.

Stefan Parkvall joined Ericsson in 1999 and is a Senior Expert working with 6G and future radio access. He is one of the key persons in the development of HSPA, LTE, and NR radio access, and has been deeply involved in 3GPP standardization for many years. Dr. Parkvall is a fellow of the Institute of Electrical and Electronics Engineers (IEEE) and is co-author of several popular books such as 4G – LTE/LTE-Advanced for Mobile Broadband and 5G NR – The Next Generation Wireless Access. Dr. Parkvall has more than 1500 patents in the area of mobile communication and holds a PhD in electrical engineering from the Royal Institute of Technology (KTH), Stockholm, Sweden.

Gunnar Mildh is a Senior Expert in radio network architecture at Ericsson Research. He received his MSc degree in electrical engineering from the Royal Institute of Technology (KTH), Stockholm, Sweden in 2000 and has since that time worked at Ericsson Research on standardization and concept development for GSM/EDGE, HSPA, LTE, and 5G NR.

Erik Dahlman joined Ericsson in 1993 and is currently a Senior Expert in radio access technologies within Ericsson Research. He has been involved in the development of wireless access technologies from early 3G, through 4G LTE, and mostly 5G NR. He is currently focusing on the evolution of 5G as well as technologies applicable beyond 5G wireless access. He is the co-author of the books 3G Evolution – HSPA and LTE for Mobile Broadband, 4G – LTE and LTE-Advanced for mobile broadband, 4G – LTE-Advanced Pro and the Road to 5G and, most recently, 5G NR – the Next Generation Wireless Access Technology. He has a PhD in telecommunication from the Royal Institute of Technology (KTH), Stockholm, Sweden.

Bipin Balakrishnan is a Senior Researcher focusing on future device technologies at Ericsson Research. He holds an MSc in electrical engineering from the Royal Institute of Technology (KTH), Stockholm, Sweden. Since graduating in 2008, he has worked on mobile device system architecture and concept development. He has also held leadership positions in MIPI Alliance, a standardization body for mobile devices. He joined Ericsson in 2019.

Peter Öhlén is a Principal Researcher at Ericsson Research Networks, focusing on service and network automation across different technology domains—transport network, radio, and cloud. The goal is to achieve full automation by efficient frameworks, exposure of relevant data, and applying smart algorithms. A key current is understanding when and how AI can be applied to deal with problems where effective solutions are not available today. Peter joined Ericsson in 2005 and has worked in a variety of technology areas in wireless and fixed networks. He holds a PhD in photonics from the Royal Institute of Technology (KTH), Stockholm, Sweden.

Elmar Trojer is a Research Leader at Ericsson Research Networks, currently focusing on split RAN architectures and fronthaul transport solutions for 5G and 6G radio networks. He holds a PhD in electrical engineering as well as an MBA and has carried out research on fixed access, small cells, 4G/5G backhaul, fronthaul, and lower-layer splits.

Göran Rune is a Principal Researcher in system and network architectures at Ericsson Research, a position he has held since 2014. He received an MSc in applied physics and electrical engineering in 1986 and a Licentiate of Engineering (Lic. Eng.) in Solid-state physics in 1989 from the Institute of Technology at Linköping University. He joined Ericsson in 1989 and has since then worked on the core network as well as radio network architecture in systems design as well as concept development and standardization for most digital cellular standards, including GSM, PDC, WCDMA, HSPA, and LTE as well as the 5G core network.

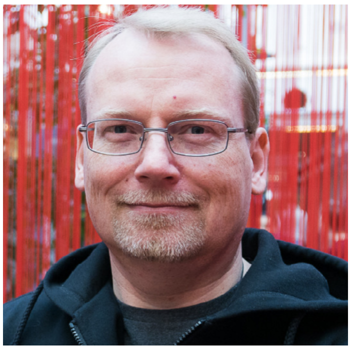

Jari Arkko is a Senior Expert in internet architecture at Ericsson Research. He has worked on software development, routers, security, mobile networking, and internet technology. Jari has also served as the Chair of the Internet Engineering Task Force (IETF).

Zoltán Turányi is an Expert in control architectures in research area cloud. His ongoing interest is in cloud-native friendly mobile architectures, function-as-a-service, and cloud execution environments. He holds an MSc in computer science, joined Ericsson in 1996 and his past interests include IP QoS and mobility, mobile core network architecture, and software-defined networking.

Dinand Roeland is a Principal Researcher at Ericsson Research. His current research interests are introducing AI technologies in the end-to-end network architecture, with the goal to achieve an autonomous cognitive network. Since joining Ericsson, he has worked in a variety of technical leadership roles including product development, concept development, prototyping, and standardization. He holds an MSc cum laude in computer architectures and intelligent systems from the University of Groningen, the Netherlands.

Bengt Sahlin leads the Security Group in the Ericsson Research NomadicLab in Finland. He joined Ericsson in 2000, initially working on mobile system security and product security. His engagement in standardization includes participating in 3GPP, ETSI, GSMA, and IETF. From 2010 to 2013 he served as the Chairman of the 3GPP security working group, TSG SA WG3.

Wolfgang John is a Principal Researcher in Network-Compute Convergence at Ericsson Research. His current research focuses on edge and distributed cloud computing concepts for both telco and IT applications. He holds a PhD in computer engineering from Chalmers University, Sweden. Since joining Ericsson in 2011, he has also carried out and lead research on network management, software-defined networking, and network function virtualization.

Joacim Halén is an expert in distributed cloud software design at Ericsson Research. He holds an MSc in engineering physics from the Royal Institute of Technology, Sweden, and joined Ericsson in 1997. He has worked on software architectures, and prototype development on all types of systems and levels. Over the last ten years, he has focused on cloud technologies.

Håkan Björkegren is a Principal Researcher who has been with Ericsson Research in Sweden since 1995, where he has focused on the air interface and radio protocol design for wireless systems. He has a PhD in signal processing from Luleå University of Technology, Sweden. He is currently working on designing and evaluating different air interface concepts for 6G.