making communications work
for everyone

Internet Futures

Spotlight on the technologies which may shape the Internet of the future

Published 20 July 2021

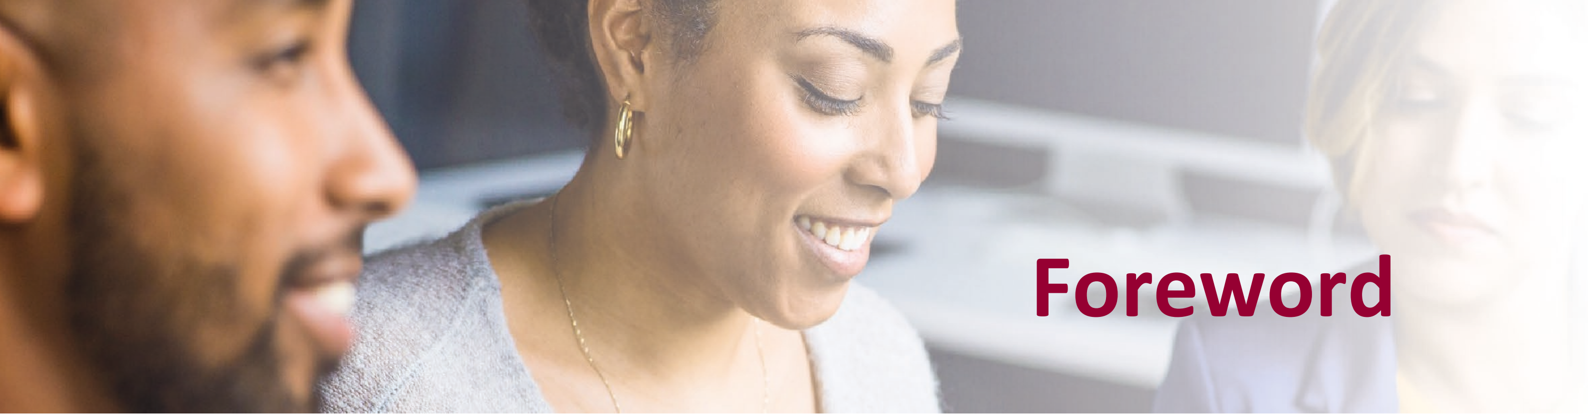

Foreword

Computers excel at producing and processing data, while telecommunications networks enable the rapid and widespread transport of data. The combination of the two gave birth to the Internet, which has been transforming the way people live their lives and do business. Today, we find the Internet has grown to an unprecedented and global scale. Many services in communication, commerce, collaboration tools, education, and entertainment only exist because of the Internet and its ability to provide a platform for innovation and reach to potentially every corner of the world.

The fundamentals of what can be computed were established by Alan Turing in Cambridge in 1936, and then turned into a practical computer architecture by John von Neumann in 1945. Networked computers communicating through the protocols established by Vint Cerf and Robert Kahn in 1974 led to the creation of the Internet, and then the invention of the World Wide Web by British engineer Tim (now Sir Tim) Berners-Lee in 1989. And ever since those early years, Internet technology has been constantly evolving and progress is being made to improve the Internet ability to transport data in terms of speed and volume and global reach. Similarly, web technology is being advanced to meet future needs for information sharing.

As the UK’s independent communications regulator, it's essential Ofcom keeps aware of changing technology. This allows us to consider how these changes can affect the sectors we regulate now, and in the future. And it informs the actions we take to make sure people and businesses in the UK continue to enjoy high-quality communications services and are protected from any risks these new technologies could pose. To this end, we maintain our awareness of the innovation in the Internet and the telecommunication industry. While many emerging technologies are already well-known to us, we recognise there will be others that are not, which could still have a major impact on the consumers of tomorrow. So in 2020, we decided to ask the world’s leading technologists and social scientists for their views on what the next game-changing technologies could be. We carried out dozens of interviews and also invited anyone with insights and evidence on new technologies to contribute to our research. We published key findings from that exercise in the telecoms space in January 2021 in our *Technology Futures* report [Ofcom2021] – this follow-up report covers our findings relating to the Internet.

We discovered a huge range of exciting technologies relating to the Internet. Some will lead to new, richer communication experiences, involving immersive technology, that enables us to touch or move objects at a distance. This could potentially increase the levels of collaboration between people, robots and things in a more converging real and virtual worlds. Others could help improve the quality of experience for consumers and businesses and facilitate information sharing with enhanced trust, privacy and security for personal and enterprise use that would improve personal access to information and sharing; advance personal communications; business collaboration tools; gaming and entertainment to name a few. Architectural advances in cloud and edge computing could enhance the efficiency and performance of many services. They may also contribute to ensuring the environmental sustainability of Internet services.

This report should not be seen as an exhaustive list of every innovative technology being developed. Indeed, it can be no more than a sample of the high-quality ongoing research work being conducted in industry and academia. Further, the omission or inclusion of any technology shouldn't be taken as a signal of our view of its importance. Nor are these our predictions for the future: this report is a summary of the technologies that have been flagged to us by worldwide experts.

But we hope that our findings offer some insight into how innovation can ensure a bright future for the Internet in the UK and the wider international community. We will continue to develop our work in this area as we engage further with people and businesses across the communications world to identify the technologies of tomorrow – and what they could mean for you and me as consumers.

Our aim remains to continue this conversation and play our part in helping the communications industry to constantly evolve and innovate.

Finally, I'd like to express our profound thanks to the many experts around the world who have shared their time and inspirational thoughts with us[1].

Prof. Simon Saunders

Director, Emerging and Online Technology

[1] See Annex B and C for a list of contributors.

Contents

Summary

Our lives have been profoundly impacted by the Internet and the web. They have made the world a more connected, but also complex place. Billions of people are now online, thus creating a very busy ‘global village’[2]. Indeed, more than 51% of the global population (or 4 billion people) were using the Internet in 2019 and exchanging traffic across the globe which has grown 12-fold since 2010, or around 30% per year) [ITU-Stats, Kamiya2020]. Many services in communication, commerce, collaboration, education and entertainment only exist because of Internet and communications innovations rolled out on a global scale. Indeed, some of the biggest companies in the world were founded on the Internet and some owe their existence to the web.

As a result, the Internet and the web have become an integral and near-essential part of the fabric of our society.

But coupled with the rise of the many benefits of these innovations, there has also been a parallel rise in concerns about equality of access, online safety, the impact of misinformation, disinformation, privacy, and security; as well as the scale and ubiquity of a number of large scale companies in areas such as communications, computation, news and advertising.

“But that global village brings in a lot of different people, a lot of different ideas, lots of different backgrounds, lots of different aspirations” [Brahimi2005]. This has played a significant role in making the Internet a fertile ground for innovation, which is played out into two ways. Firstly, it helps advance the Internet itself as a medium for constantly improving global connectivity in terms of providing the basis for better quality-of-experience (QoE) and global reach. Secondly, the Internet plays an important role as a platform for innovation which depends on its ability to provide QoE with global reach with no inherent limitations on who can access, build, and study it.

This leads to what has been termed a ‘permission-less’ [IS2019] form of allowing anyone with the relevant skills to create a new, globally available service as long as they conform to the existing interoperable Internet standards and best practices. This is in addition to the role played by open source software in lowering the barriers to entry for providers and vendors.

The importance of the Internet and its impact on the world stems from the fact that it is sandwiched between the people and content and application providers, see Figure 1, combining two functions, namely data transfer and data compute [O’Hara&Hall2021].

[2] The term ‘global village’ was coined by Marshall McLuhan in 1962 [McLuhan1962].

Figure 1: Overview of today's Internet. Source: Ofcom.

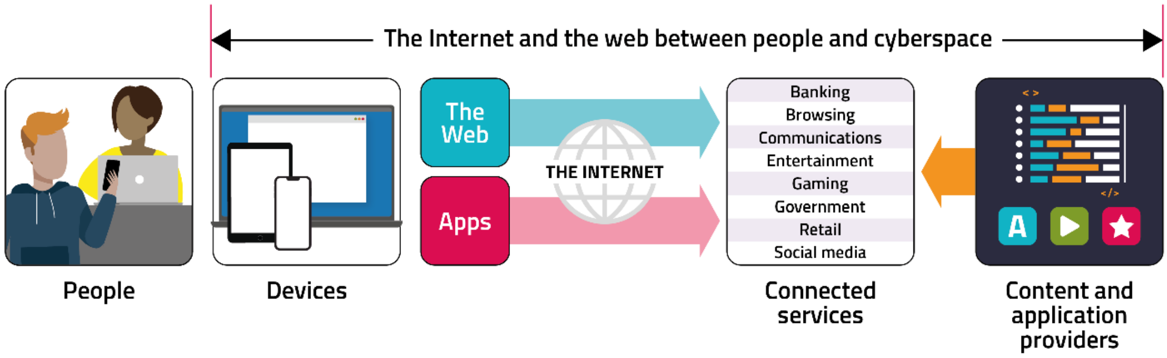

Scope of our emerging technology work

Internet technology continues to evolve across all of the industries Ofcom regulates – bringing new products, services and ways of working for people and businesses. This report focuses on innovation in the Internet sector. Its aim is to increase our understanding of the way technology is changing, so we can help ensure these advancements benefit people and ensure they are protected from harm.

Figure 2: Scope of Ofcom's emerging technology programme. Source: Ofcom.

As the UK's independent communications regulator, our work is wide-ranging – covering the Internet, broadband and mobile; TV; radio; video-on-demand services; postal service; and the airwaves used by wireless devices. The UK Government introduced new legislation in autumn 2020 giving Ofcom powers to regulate UK-based video-sharing platforms. The UK Government has also confirmed Ofcom's appointment as the regulator for online safety. We are also working with the Government to implement new telecoms security legislation.

To provide support for our work across this diverse range of sectors, we have developed our emerging technology programme, see Figure 2. The programme looks at how the way we communicate and connect with each other is evolving, both across the Internet and the telecoms networks that carry it.

Scope of this report

In January 2021, we published our Technology Futures report which examined the innovative emerging technologies that could shape the telecoms industry in the future [Ofcom2021]. This report is a follow-up to identify and illuminate developments that could shape the future Internet.

This report is a summary of what we learnt from a series of interviews we conducted with internationally renowned subject matter experts including: Internet pioneers, futurists, standardisation experts, academics, social scientists, and people in the industry. We also incorporated the views of people and organisations expressed in a public call-for-input and drew from our own supporting in-house desk research. See methodology section.

While we are generally interested in the implications that longer-term Internet technology can have for people and society, we do not assess or describe in this report the policy implications of the technologies we spotlight, including service consistency, reliability, privacy and security; consumer harm; and future legislation and regulation.

Identifying new technologies

During our work on emerging technology, we have looked at how existing Internet technologies are evolving, and those that are starting to emerge now – and in the coming years. In particular, we have focused on identifying Internet technologies with the potential to disrupt the Internet space in the medium to long term. See methodology section for further details.

Emerging trends

We found a number of developments which might shape the Internet of the future. The majority of what we found revolves around key themes:

1) Personalisation of experience, privacy, trust, and QoE;
2) Increased introduction of cloud technology and its extension closer to the end user; and
3) Awareness of environmental sustainability.

Figure 3 gives further details of the developments that underpin the above. The vertical lines represent the technological developments that impact the quality-of-service (QoS) in terms of effective, efficient and resilient application traffic delivery. The horizontal lines depict the trends that are associated with future personal and business applications in terms of QoE. These reflect the current direction of travel of applications and services to be delivered using future traffic delivery technology represented by the vertical lines.

Figure 3: Technology trends across the Internet sector. Source: Ofcom.

We next highlight the key emerging trends from the four chapters contained in this report.

Everyday life with the future Internet

As Internet technology continues to evolve, a key future trend that emerges is the concept of the Metaverse[3]. Understood as virtual shared space that converges with actual reality, the Metaverse may allow people to come together to work and socialise without having to log on to different devices or platforms. In theory, this continuum would allow us to seamlessly transition between the daily activities that, by this time, will be conducted online. The prevalence of the Metaverse in gaming might be one way in which it spills over into our everyday life.

One direction of Internet evolution is to support sensory interactions, where people could interact with the Internet through their senses such as touch, smell and taste, in addition to sight and hearing: known as the Internet of Senses [Ericsson2019]. This would enable greater personalisation of individuals' experiences. The conjunction of virtual reality, augmented reality, and further developments in technology such as artificial intelligence, might lead to more enriching immersive experiences. This includes applications ranging from gaming and entertainment to retail, work and education. If the Internet of Senses becomes a reality, the technology could evolve to allow people to meet and interact in ways similar to face-to-face and as natural as possible. In this world, the boundaries between reality and virtuality could become increasingly blurred, as the interaction with the Internet is likely to become less mediated by screens [DeNardis2020].

Artificial intelligence and the Internet of Things underpin smart cities. In such cities, people could become more of 'passive agents' through being connected via intelligent devices (e.g., wearables, chip implants, etc), meaning that people may be less aware of where and when their interaction with the Internet stops and/or begins.

One direction of Internet evolution within industry is to support increasing levels of seamless working between humans and robots in both real and virtual worlds, which could further blur the boundaries between the two.

Interoperability and digital identity are already important and could become more so in an ever- evolving connected world along with the notion of greater trust and privacy.

Evolution in the Internet and web technology

The Internet standards continue to evolve with backward compatibility which is key to maintaining continued interoperability with legacy protocols and devices. There is significant innovation taking place in the transport layer be it from an optimisation perspective. The standardisation of the QUIC protocol is key a milestone which is expected to improve traffic delivery and ultimately QoE, with further enhancements are underway to QUIC and congestion control.

While we see that the IP version 6 (IPv6) adoption is on the rise, there is also ongoing effort to advance and enhance the future capabilities of IPv6 – this is in addition to the development of interoperable alternatives to IP. Similarly, effort is also ongoing into niche areas such as delivering IP over wireless deterministic networks.

Traditionally, applications and networks have been decoupled. We are seeing recent effort to

[3] The Metaverse is a term coined by Neal Stephenson's in 1992. It is created by the convergence of virtually enhanced physical reality and physically persistent virtual space [Stephenson1992].

increase the Internet applications awareness of the network in terms of the quality-of-service available. Similar effort is being made to enable network awareness of the specifics of the Internet applications needs for resource. This may pave the way for greater levels of cooperation between networks and applications, helping applications perform more effectively and lead to a better QoE.

Routing resilience is being addressed with the ultimate goal of assuring end-to-end route security. Increased routing resilience may add protection to the Internet as a global infrastructure with continued rising levels of importance to everyday life.

The advances at the transport layer are also being reflected in web protocols with the rise of HTTP/3, leading to improved web-browsing experience. Further, we also saw developments towards the semantic web which caters for trust and privacy through aspects such as improvements in data ownership, for both personal consumer use and enterprise. This could pave the way for trusted information sharing such that only the necessary information is actually shared without compromising consumer expectations and business needs for data.

Evolution in the Internet compute, connect and cloud

As Moore's law starts to wane, interest in alternatives to the successful von Neumann processor architecture would continue to rise. Emerging computing architectures like neuromorphic computing, memory-centric computing, optical computing and quantum computing continue the architectural innovation to increase growth in compute capacity. Even if the traditional von Neumann architecture remains the foundational concept, the innovations in computer hardware architectures continue to develop in parallel with innovations in the software development area.

We are seeing evolution in geographically distributed compute, storage and network technologies in the form of cloud computing. Further, the edge of the network and the traditional cloud appear to be merging so that all cloud infrastructure within the large regional data centres to the smallest edge can be viewed as a single, unified, integrated execution environment for distributed applications – this is sometimes referred to as the 'cloud-edge computing continuum'. It offers geographically distributed compute and storage resource with variations in the amount of capacity it can serve and the level of latency it can meet. As a generic computing resource in nature, it may be used to serve both end applications as well as network functions within the wider context of virtualisation for consumer, business, and network operation applications.

Similarly, virtualisation technology is evolving to offer cloud operators and applications developers a wider range of more efficient and diversified set of development tools.

The Internet and energy efficiency

While it is important to note that efforts are being driven globally to reduce the carbon emissions across the entirety of the information and communications technology sector, including in areas such as consumer devices and network equipment, our focus in this report is on the opportunities to reduce energy consumption in tomorrow's Internet cloud and edge infrastructure, spanning across data centre, mobile network, and fixed network operators [Ofcom2021].

We are seeing continued effort being made to increase the opportunities to reach the target of net zero carbon emissions by 2050 by diversifying the toolbox available for network operators. For example, we learnt that the opportunities to reduce energy consumption in future cloud-edge infrastructure could come from the basic building technology, where we see advances to reduce energy consumption by the computer processor as well as improve energy storage technology.

Next, the way data centres are built impacts energy consumption. For example, their proximity to energy sources helps lower energy wastage from source to point of consumption. The choice data centre networking of network topology also contributes to the energy consumption with mesh consumes least energy followed by star topology. Finally, disaggregated server farms within the data centres themselves also help cut down energy consumption.

From an operational perspective, dynamic real-time management of compute and storage resources with respect to load and required QoS is an important mechanism to reduce energy consumption in data centres and across the cloud-edge infrastructure.

It follows from the above that various practical network design and operation techniques are available for reducing energy consumption such as reducing the number of traffic hops, reducing the number of conversions between electrical and optical connectivity and distribution of points of service. These techniques enable dynamic management of the network in real time to provide the required quality of service with reduced energy consumption.

Looking at the Internet of tomorrow

Figure 4 provides an emergent picture of tomorrow's Internet from this research. It describes five paths for development of the Internet [Hernández-MuñozetaletalEtAl2011]:

1. The Internet for delivery of traffic, which is what we have today but with increased support for QoE and greater global reach.
2. Internet of People (IoP) enabled by seamless connection, interaction and information exchange by people about within their social context and environment.
3. Internet of Things (IoT) which is a network infrastructure where physical and virtual 'things' are seamlessly integrated into a common information network through which interaction takes place[4].
4. The Internet of Services (IoSr) is created by various interoperable applications through the use of semantics for the manipulation of data and information from different service providers, sources and formats.
5. The Internet of Senses (IoS) which enables people to use their senses to communicate with each other as well as with machines and things over the Internet.

People, machines and things interact with each other in the IoP, IoT and IoSr. The Internet of senses could facilitate this interaction through a combination of future evolution of the web and devices, and with trust. The Internet for delivery continues to act as the 'workhorse' that delivers traffic efficiently amongst all other constituents with the cloud-edge

[4] IoT devices have become widely adopted but the heterogeneity in devices' user interfaces, protocols, and functionality still present a challenge for seamless interoperability and use [RenEtAl2019].

infrastructure providing the requirements for compute and storage in an environmentally sustainable fashion.

Figure 4: Potential Internet of tomorrow based on this study. Source: Ofcom.

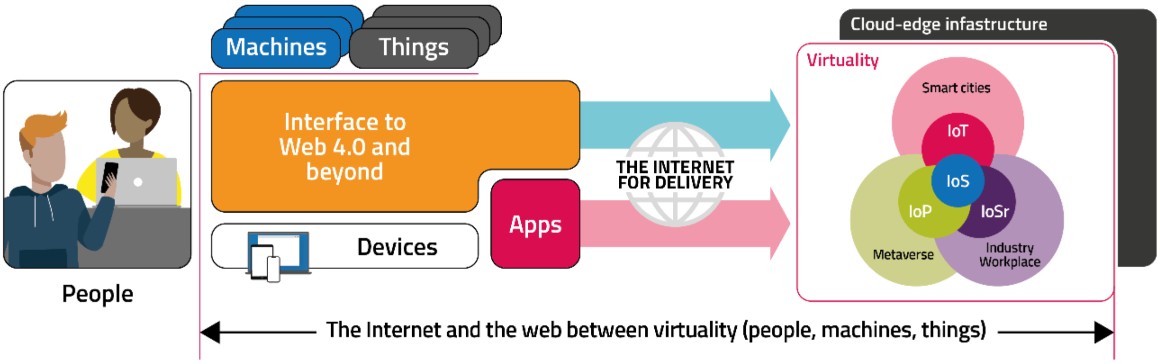

Methodology

The content of this report arises from our need to understand the potential role of emerging technologies in the sectors we regulate, and to guide our work programme towards the technologies most likely to have the greatest impact. To this end, our objective was set to capture as many Internet technologies as possible which could potentially be impactful on our sectors and consumers. We sought to minimise the chances of missing potentially important technologies – i.e. we preferred to have 'false positives' rather than miss significant developments.

We acknowledge that the future of the Internet is influenced by combinations of the geo- cultural, economic, business, social and political aspects of the world. That said, the Internet sector is also very much affected by technology innovation, which is evolving rapidly across all the industries bringing new products, services, and ways of working for people and businesses.

We acknowledge that the technologies reported are at various stages of evolution, some of which are in a more advanced stage in development than others. Consequently, some of the technologies may emerge in the next 5 years, whilst others may take much longer to emerge.

Thus, we are primarily interested in technologies which have the potential to significantly impact on at least one, and typically several, of the following criteria:

- Enables the delivery of new services which are valued highly by people and businesses.
- Increases the performance of networks, improving the experience for people.
- Lowers barriers to entry for providers, enabling choice for people.
- Reduces the cost of delivering services, increasing access and maximising value for customers.
- Reduces the total environmental impact of delivery of communication services and associated activities.
- Assures the security and resilience of service delivery.

The approach

We recognise that the Internet is the 'interface' between people, machines/things and cyberspace. This led us to tailor our approach to combine aspects of the technology and people dimensions (use cases).

In particular, the framework shown in Figure 5 was developed drawing from the delivery-focused Internet layered concept[5], which we complemented with the consumer at its heart, Internet access and services layers to cover our

[5] The Internet protocol design consists of 4 layers namely, the application, transport, network and network interface and hardware [ParzialeEtAl2006].

needs to explore the technology, people and machine/things dimensions. We also took into consideration the Open System Interconnect (OSI) model [ParzialeEtAl2006][6].

Figure 5: Framework encompassing all aspects of Internet within our scope. Source: Ofcom.

- 'Internet services' captures future services the end users will engage with in the future.
- 'Information/service management' captures the technology development in the area of information/service

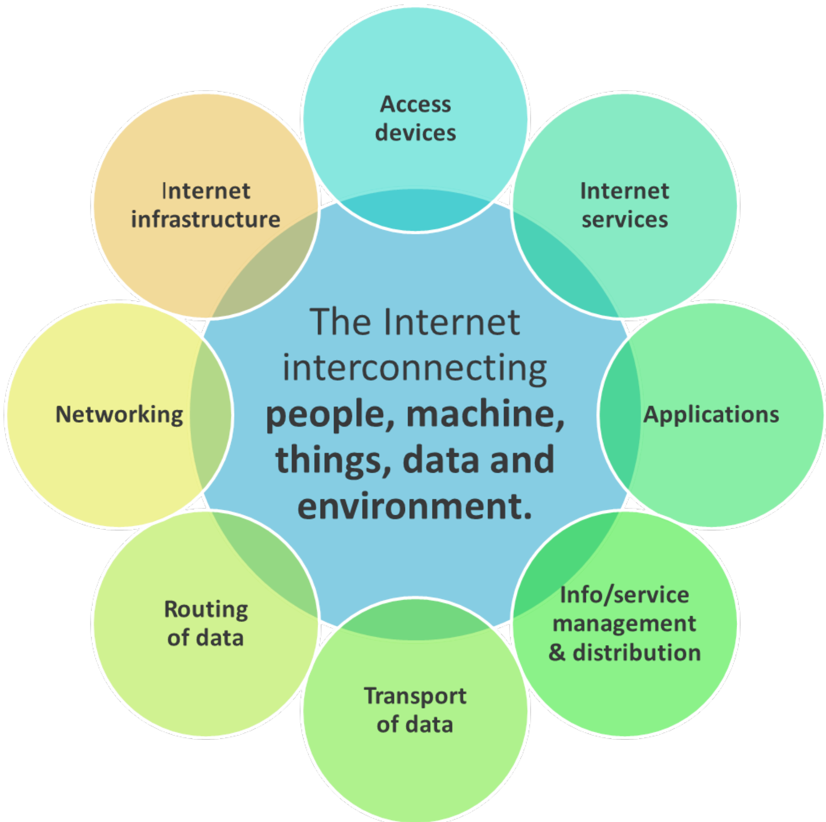

distribution, processing and management.

The framework describes the various constituents, as defined below, of the Internet space and thus constitutes our research scope:

- The end user which includes entities that produce and/or consume information/ services, such as consumers/people, machine and things, data and environment.
- 'Access devices' captures future technologies and devices the end user will use to access the Internet.
- 'Applications' captures how the end users will use the Internet (use cases).
- 'The transport of data' captures developments in the area of reliable and ordered transmission and reception of packets.
- 'The routing of data' captures developments in the area of routing technologies, which is responsible for choosing 'the best route' to deliver the data packets from the source to the destination with route security included.
- 'The network' captures developments in the IP layer which defines the addressing scheme at the network and node levels

[6] Referring to footnote (5), the mapping of the 4 layers of the Internet suite of protocols to the OSI model leads to the layers being sometimes referred to as layer 7, 4, 3 and an aggregation of 1 and 2, respectively [ParzialeEtAl2006].

and effectively specifying how packets are routed in the network.

- 'Internet infrastructure' includes future developments in the physical layer which define the electrical, optical and wireless specifications of the data connection. It captures the developments in technologies that will process the data along the path from source to destination.

Source of material

To achieve the objective for this research, we have engaged with academia and industry internationally leading on research in the fields of technology and social sciences.

We interviewed over 65 internationally renowned subject matter experts: Internet pioneers, futurists, standardisation experts, academics, and people in the industry. We ensured that the range of experts covers the framework of Figure 1, see Annex B for the list of people we spoke to in this context. We also sought the views of people and organisations through a public call-for-input and received 30 responses. In Annex C we provide the full list of responders, and a link to where the non-confidential responses can be found. We also drew from our own supporting desk research [Ofcom2020].

Structure of the report

We have split this report into four chapters, each of which is focussed on a different area of the Internet, namely:

- *1: Everyday life with the future Internet* which addresses future Internet use cases.
- *2: Evolution in the Internet and web technology* which reports on the latest developments in the Internet network (layer 3) and transport (layer 4) technologies; and above.
- *3: Evolution in Internet compute, connect and cloud*, which summarises some of the development in the processor technology, optical connectivity and cloud-edge computing.
- *4: The Internet and energy efficiency* which addresses how future cloud and edge infrastructure may be built and operated with environmental sustainability in mind.

We have mapped our framework of Figure 5 to the chapters in the following way:

- Chapter 1: covering people, machine, things, access to the Internet, Internet services and applications.
- Chapter 2: covering secure information management and distribution with privacy; transport of data; and networking and routing.
- Chapter 3: covering the Internet infrastructure in terms of compute, connect and cloud.
- Chapter 4: covering the environmental sustainability aspects.

Looking ahead

We continue our engagement with experts in other technology areas that are relevant to us. We aim to share the outcomes of this evaluation with stakeholders through appropriate publications, the exact mode and timing of which will be decided in due course. Meanwhile, we welcome views on the technologies we present here, or on technologies which we have not yet taken account of, via our email address:

emerging.technology@ofcom.org.uk

1: Everyday life with the future Internet

This chapter provides a summary of our research into the potential evolution of the Internet from the perspective of consumers, citizens, businesses, and industry. It highlights how the Internet as we know it today will no longer be serving just as a communication network but is likely to transform how we experience life[7]. The chapter revolves around three spheres that underpin our lives, namely personal communications, living space, and workplace and industry.

Context

As authors Professor Dame Wendy Hall and Dr. Kieron O'Hara observe in their forthcoming book, "'Internet' is short for 'Internetwork', and implies it is a global network of connected computer networks." [O'Hara&Hall2021]. These networks are connected by wire (such as copper or fibre optic) and wireless communications technology (radio and light) which can carry data. Layered on top of these networks, we find the rules that make the network connect, known as the Internet suite of protocols[8]. While layer 3[9] defines how computers can exchange data with each other, layer 4 governs the connections between computers to exchange data using IP.

The web, an application that runs on top of the Internet[10], was invented in 1989 by Sir Tim Berners-Lee who wanted to see scientists exchange information about their research with peers in universities and institutes around the world [Getting2007]. To this day, the Hypertext Transfer Protocol (HTTP) and the Uniform Resource Locator (URL) are two of the web's fundamental technologies. Whereas the former governs how documents get transferred across the web and over the Internet, the latter identifies a web resource and indicates its location[i] [O'Hara&Hall2021].

The success of both the Internet and the web is demonstrated by the fact that over half of the global population is online today, making the Internet an integral part of the social and economic fabric of many communities around the world, see Figure 6. This has also led to the blurring of the demarcation between the Internet and web, where people sometimes use the two terms interchangeably.

---

[7] Worth noting that the Covid-19 pandemic has already led to a major shift in the way we interact with the Internet. Daily traffic surged as people stayed at home during lockdowns and applications for remote working and education, such as VPN and video conferencing, experienced traffic increases beyond 200% [FeldmannEtAl2021].

[8] Refer to footnote (6).

[9] See footnote (6).

[10] Email is another example of a popular application that runs on top of the Internet.

Figure 6: The relationship between the Internet as a delivery medium and applications. Source: Ofcom.

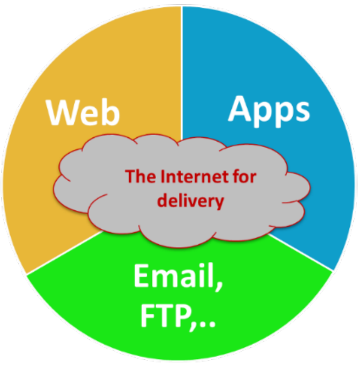

Web evolution

The web was not always the one we know and use today; there have been different versions of it since its development in the late 80s, as Figure 7 demonstrates [FWB2018].

Figure 7: Web evolution[11]. Source: [FWB2018].

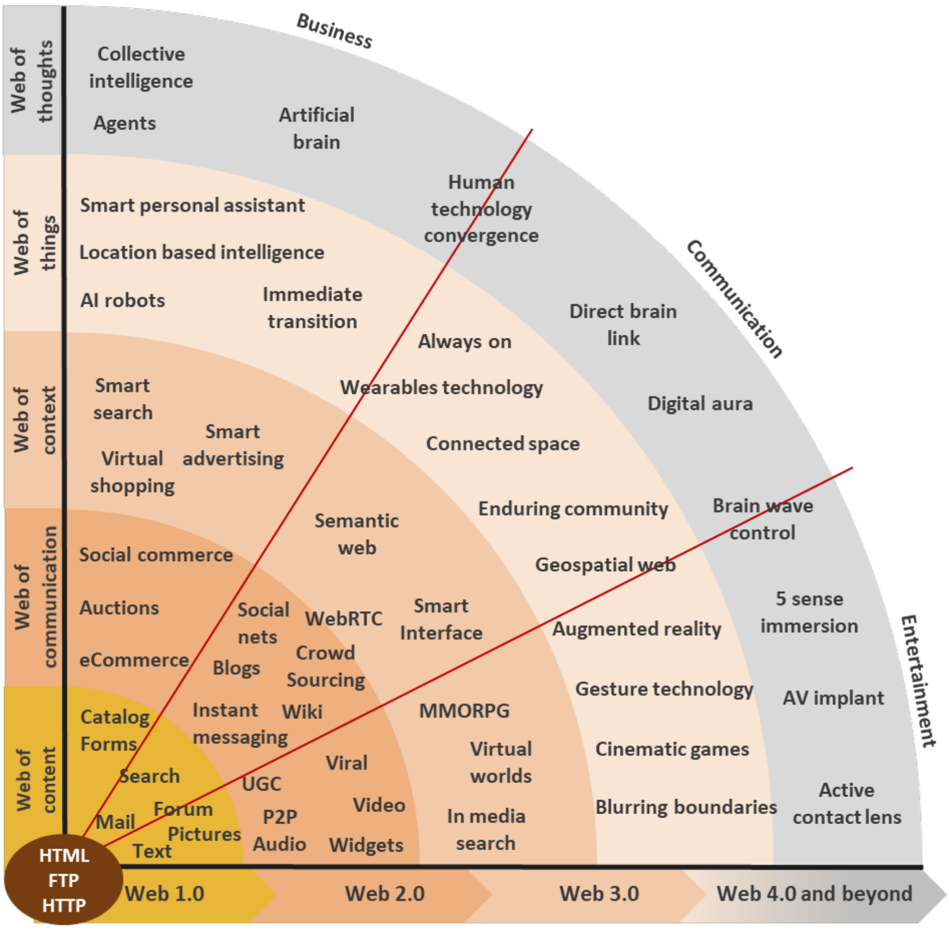

[11] © Reproduced with permission from Trendone.

Web 1.0 was initially designed as a means of information sharing between scientific institutions and universities. Often referred to as the 'read-only' web, it offered users limited to no interaction with online content [AghaeiEtAl2012].

Static websites and adverts were designed to be brochure-like in their presentation of information and products to users [FWB2018].

This began to change with the development of Web 2.0, the 'read-write'. Web 2.0 enabled consumers to publish as well as read content. The beginning of this era was defined by websites such as LiveJournal and Blogger, which allowed for online interaction between users with limited multimedia capability.

Initially limited in its interconnectivity and information sharing between platforms [Choudhury2014], Web 2.0 has grown to include more sophisticated user-generated multimedia content, interactive web applications, and social media [Rene2019].

Major Internet experts suggest that Web 3.0 will be dominated by the 'semantic web' [NathEtAl2015][12], which is depicted in Figure 7. The term semantic is used because through the analysis and structuring of data, computers will at this stage be able to 'read' and analyse web content. For the consumer, this would mean things like more refined search engines and greater personalisation of experience [Choudhury2014].

We are on the cusp of moving from Web 2.0 to Web 3.0 given that aspects of Web 3.0 are starting to emerge, such as the semantic web implementations and greater levels of personalisation.

Through innovations in virtual reality, augmented reality and artificial intelligence (AI), the information contents of the Internet will become seamlessly interlinked with the physical world [CookEtAl2020]. With this innovation, the clear distinction we see today between the physical and digital worlds starts to dissolve.

Another defining trait of Web 3.0 is its mobility [PandoraFMS2018]. In today's world, many of us access the Internet through our mobile phones as much as our computers. This, combined with cloud technologies, means that Web 3.0 could be the most accessible iteration of the web so far.

Looking to the future, Web 4.0 and beyond has been under discussion [Almeida2017]. Our research suggests this future web iteration could combine the real and virtual worlds in a seamless fashion. The services provided could then be autonomous, proactive, content-exploring, self-learning, collaborative, and use content-generating agents based on fully matured semantic and reasoning technologies as well as AI [Weber&Rech2009].

Some researchers have pointed out that from a consumer perspective, the future is the sensory and emotive Web 4.0, where machines interact with human beings in a much more collaborative fashion [AghaeiEtAl2012]. This relationship could become the norm for many people. Currently, the web is emotionally neutral in the sense that it does not perceive what users feel. AI algorithms are advancing to become attuned to human feelings and emotions, thus vastly improving interaction with devices[13]. This vision for the future web would enable consumers to interact with content that responds to their emotions. These highly personalised interactions enable the consumer experience of this future web generation to be more affable than its predecessors.

The Metaverse

'Cyberspace' as a term was first coined by William Gibson in his 1984 book, Neuromancer [Gibson1984]. A cyberspace is a digital, virtual world in which users share information, interact, swap ideas, play games, engage in discussions or social forums, conduct business, and create intuitive media, among many other activities. Up until now, these cyberspaces

[12] Web 3.0 has also been referred to as the 'spatial web' [CookEtAl2020]

[13] Emotion detection technology consists of two components: computer vision, to precisely identify facial expressions, and machine learning algorithms to analyse and interpret the emotional content of those facial features [Schwartz2019].

have been largely separated from one another, with different spaces for different activities.

The Metaverse is the concept of a virtual shared space that converges with actual reality [Stephenson1992]. The idea is that all of these different real and virtual spaces merge into one continuum where people will be able to work, socialise, collaborate and interact with one another without any delineation, just as they do in the physical world. Further, a consumer would be able to move seamlessly between virtuality and reality thus interacting with other real or virtual people in both worlds.

The Metaverse concept is currently most prevalent in gaming [Stephenson1992]; think massive multiplayer online games such as World of Warcraft, where an online world comprised of other gamers exists regardless of whether you yourself are online.

The prevalence of the Metaverse in gaming might be one way in which it spills over into our everyday life. For example, after finishing work for the day, be it in real or virtual space, people who want to game with friends might no longer need to log onto another platform but instead find their friends in another corner of the Metaverse.

The same technology has been used to host multi-room collaborative video/audio meetings, an area of experimentation that has grown in the past year and could mimic real world experiences [Kwok2021].

The Metaverse may have a significant effect on future Internet use. The creation of this virtual space would be a new and exciting development, possibly serving as the next new platform for media content creation; think along the lines of a virtual Travis Scott performing a concert in popular online game Fortnite [Webster2020]. If the Metaverse continues to develop from its current, nascent stage, its influence could be far-reaching.

The Internet of senses

The Internet of Senses (IoS) could underpin the potential shift from the emotionally neutral web to a multisensory experience in cyberspace. While still far from reality, the concept of extending our sensorial experiences to the Internet has been considered for years, as part of on-going research and development projects. It is not surprising that the arrival of the IoS has been identified as one of the most awaited trends over the coming decade [Ericsson2019][14].

The ‘raw technologies’, which could facilitate an immersive experience, such as those reviewed in [Ofcom2021] represent some of the key enablers of the IoS, all of which will bring about the multi-sensorial Metaverse where people can interact using their senses[15].

If this concept materialises, the IoS might change our screen-based interactions with the Internet and could become a component of the future of the web. In this way, the IoS adds new components to the current human-machine interface devices (keyboards, mouse, and game consoles) as people start to use gesture controls and potentially brain activities directly to initiate commands. For example, the

[14] Some may remember Google’s Nose Beta in 2013, a search engine that would enable users to submit search queries not by text, but smell. The announcement was in fact an April Fool’s joke, but one can see how this novel concept creates enthusiasm [Google, Castro2021].

[15] An example of such enabler is the use of haptic technology which recreates the sense of touch and is already widely used in mobile devices and game controllers [McClelland2017]. Similarly, haptic technology has been applied to Braille perception using a touchscreen and vibration on mobile phones [JayantEtAl2010].

adoption of augmented reality glasses, could mean that by uttering the words or thinking 'show map', people would see a map display before their eyes. This in turn would allow them to search for routes for example, simply by using voice commands or thinking of their journey destination.

Consequently, as the IoS develops further and consumer engagement increases, the Metaverse may blend with our own everyday physical reality, reaching a point where there could be no distinction between the real and the virtual worlds.

Future advances in the IoS might also instigate a significant impact on how people will work in the future [Ericsson2020]. The pandemic has already ushered in a more flexible and home-based working pattern. As such, more immersive and sensorially rich experiences could become part of our professional lives, as advances are made in the relevant key technologies. This could include, for example, meetings in the Metaverse that replicate actual work environments rather than just viewing participants remotely; or a wearable device that uses digital technology to mimic the temperature in an office or conference room, enabling employees to feel cool even when it is hot at home [Ericsson2020].

The ability to integrate more sensorial experience into the Internet could also transform online retailer and shopping habits. Research has shown that some people need to touch the product (or imagine touching it) to be confident about their choices [PetitEtAl2019]. The IoS might enable users to immersively interact with the products to confirm its dimensions, texture, and maybe even weight.

Additionally, the web archives themselves have been recognised as sources of historical information [Brügger&Schroeder2017]. There could be a possibility in which the increasing ability to digitise historical documents and manuscripts through advances in handwriting recognition software might potentially extend the reach of the IoS into history, thus extending the reach of our sensorial experiences into the past, see Figure 8.

Figure 8: AI-aided extended reach of Internet of Senses. Source: Ofcom.

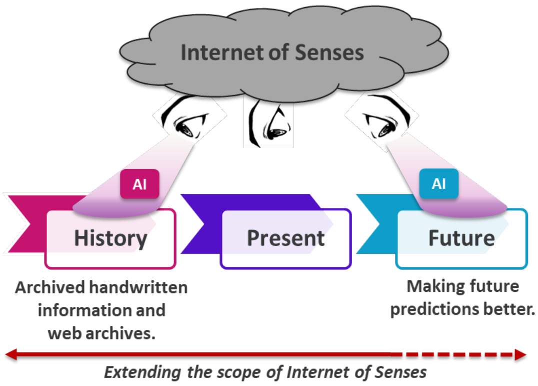

The examples above further highlight that the delineation between real and virtual space might become more blurred and no longer be mediated by screens and keyboards alone [DeNardis2020].

Digital twins, MetaHumans, avatars, and robots

What might people look like in the Metaverse? As gaming technology advances and the appearance of the characters in these games become more lifelike, so too could the appearance of digital entities in the Metaverse.

A digital twin is a model of a real entity [Saracco2019]. They represent precise virtual copies of real humans, machines or systems and are based on computer models that mirror almost every facet of the 'thing' or of the human [TaoEtAl2019, Tao&Qi2019, Guo2021].

MetaHumans are similar to the human digital twin or virtual human avatars. They are photorealistic digital avatars that exist within the Metaverse [Statt2021]. A MetaHuman is constructed from detailed images of a person to produce a computer graphical model of the respective person. The virtual avatar's movement can be synced to the movement real human thus mimicking the latter's every expression, posture and physical gesture. This type of approach could give humans a representation in the Metaverse [HemphillEtAl2019].

These avatars can be digital stand-ins and used, say, for training and education purposes, or even how we choose to appear to other people in the virtual space. In the real world, they are robots.

Figure 9 describes the different forms of MetaHumans in the Metaverse.

Figure 9:. MetaHumans in the Metaverse. Source: Ofcom.

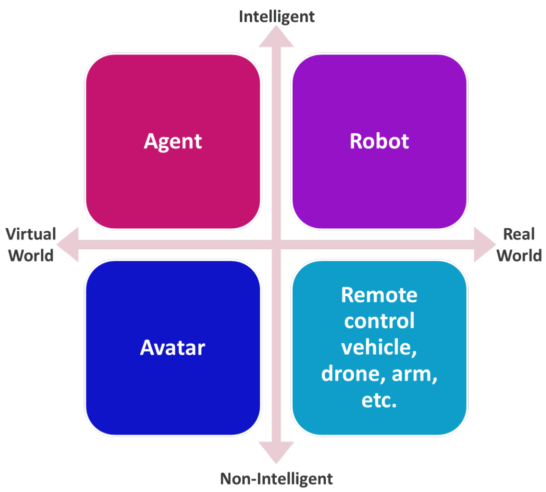

Digital personal assistant

Digital personal assistants are technically referred to as intelligent agents or sometimes just agents. In effect, they are the digital equivalent of an assistant who is able to autonomously complete certain tasks for an individual.

We can see the beginnings of intelligent agents in smart speakers such as Amazon's Alexa or Google Nest[16]. These software-based devices are capable of receiving input (in this case voice), understanding commands, and then executing the tasks that corresponds to the command, for example, setting a calendar reminder. Intelligent agents' uses would be numerous, from personal to business, such as customer service or administration.

Digital personal assistant technology has its roots in the concept of agents (alternatively known as software robots) which has been developing in the fields of computer science and AI [Wooldridge&Jennings1995]. An agent is a software-based system (entity) that can operate without the direct intervention of humans within the Metaverse, as a software robot in virtuality and as a robot in the real world, with the ability to move in between. They are autonomous entities that may be dispatched through the Metaverse to perform transactions and retrieve information on behalf of their owner [Wooldridge&Jennings1995][17].

Some researchers claim that future agents might be equipped with human-like behaviours and this is an on-going area of research [CBS2021, Moore2018, CornellU2018, Thilmany2007]. Such human-agent/robot interaction could be a factor in achieving greater integration into human lives.

Agents/robots are said to need the capability to understand human natural language; recognising human emotions; and give a response that is both understandable and

[16] Intelligent agents as a standalone product. Some might view voice assistants in mobile devices as their precursors.

[17] Web technology has been shown as one potential means of achieving primitive forms of agents, [LaclavikEtAl2006].

interpretable by humans [Gujral2019][18]. ‘Artificial emotional intelligence’ could allow agents to detect people’s emotions from the subtle cues in their voices or faces, or both, and respond accordingly. Another characterising feature could be that agent/robots do not communicate false information, nor have conflicting goals, and hence will always try to do what is asked of it [Wooldridge&Jennings1995].

Smart cities

As highlighted earlier in this chapter, a potential feature of the future of the Internet might be a world in which screens play a smaller role in mediating our online interactions and our ability to be ‘connected’ would transcend computers, mobile devices, wearables and gadgets, to connected spaces constantly gathering and transmitting information via a myriad of connectivity protocols, platforms and technologies [Baker2015]. We are already seeing the emergence of this screen-less trend with voice enabled technology increasingly taking a prominent role in the way families engage at home [Hall2020]. Beyond the home environment, we are seeing this trend at the macro level, with the rise of smart cities: urban or community spaces where existing networks and infrastructure are integrated with various digital technologies designed to benefit the residents and businesses[19]. At the core, the Internet plays, and will continue to play, an integral role in the existence and evolution of these connected environments.

Smart cities are underpinned by the growth and autonomous capabilities of Internet of Things (IoT) technologies, objects embedded with sensors that can communicate and share data with people and other connected devices over the Internet. Through these sensors, cities can gather enormous amounts of data that can be quickly analysed and made to provide near to real time information and predictions. Globally, we are already seeing examples of smart city applications with IoT solutions developed to address common needs of urban and the wider community environments. These include Singapore, where a real-time traffic management system has been put in place to help manage traffic congestion [LandTransportAuthority], and Barcelona, where in some areas waste is disposed in a subterranean vacuum network, also referred to as a pneumatic system, which includes bins equipped with sensors that monitor waste levels and help the city optimise waste disposal collection times [AjuntamentDeBarcelona].

As we look into the future, ‘vertical’ IoT applications that address specific challenges could seamlessly integrate into a common information network through open, flexible and standardised protocols and interfaces. This integration could bring together the various services into an interoperable global platform(s), also described as ‘Internet of Services’ [Hernández-MuñozetaletalEtAl2011]. The growing capacity of cloud computing will

[18] This technology is in use today in for a variety of purposes, including chatbots, market research and more [Gujral2019].

[19] It can also help guide policy-making decisions, as seen during the Covid-19 pandemic. The Royal Society’s Rapid Assistance in Modelling the Pandemic (RAMP) used urban analytics and spatial modelling to model the transmission of the virus in urban areas and small spaces which helped policy makers understand the different options for interventions and their likely consequences in the spread of the virus [TheRoyalSociety2020]. Urban analytics draws from data which are captured by governments, businesses and other intermediaries, which form the basis of Smart Cities.

continue to be an important enabler of this evolution as governments seek to grow smart cities into large scale, ubiquitously connected urban environments, ultimately envisioned as "ecosystems of embedded spatial intelligence" [Komninos2013]. Being part of this ecosystem, people could have the ability to seamlessly tap into these intelligent networks, connecting and sharing information with others in the network, and thus forming the 'Internet of People' as another pillar of smart cities [Hernández-MuñozetaletalEtAl2011].

Table 1 lists a number of key technologies and components for smart cities [SchaffersEtAl2011]. The role of the Internet is clearly evident from the table in terms of providing the delivery infrastructure.

Table 1: Technologies and components for Smart Cities. Source [Hernández-MuñozetaletalEtAl2011].

	Technologies
Content management tools	Media Internet technologies; Scalable multimedia compression and transmission; Immersive multimedia.
Collaboration tools	Crowd-based location content; augmented reality tools; Content and context fusion technologies; Intelligent content objects; large scale ontologies and semantic content.

Finally, it is worth mentioning that the concept of 'smart villages' has been gaining traction in recent years and its importance might grow [ENRD2018]. Smart villages primarily rely on a participatory approach, meaning involvement of local communities, in developing and implementing strategies to improve their economic, social and/or environmental conditions. In such settings, the use of digital technologies can be a tool in addressing local challenges and achieving rural development goals [ENRD2018].

Examples of such already existing applications include smart irrigations systems, enabled by sensors, reach of telemedicine with remote diagnostic and screening devices and use of energy efficient solar powered products [WB2021]. The increasing emphasis on sustainable development goals could lead to more efforts in bringing the benefits of applicable smart cities solutions to rural contexts especially as roll-out and accessibility of high-speed broadband, accompanied by digital literacy training, ramp up.

Industry 5.0

Industry 4.0 is where we currently find ourselves in a process that began with the first Industrial Revolution[20]. Industry 3.0 was marked by the introduction of computers to industrial manufacturing and production. However, what ultimately drove these automated computer processes was human input, not data. The defining trait of Industry 4.0 has been the shift from human input to data as the driver of automation [OUK2019]. This data is generated and utilised in many different ways, from the interconnected, intelligent devices of IoT, to cloud computing and cyber-physical systems. The more data these smart machines consume, the more

[20] The first Industrial Revolution occurred towards the end of the 18th century and saw the introduction of mechanical production thanks to water and steam power. The second Industrial Revolution followed at the beginning of the 20th century, when electrical energy gave rise to mass production [Deloitte2015].

efficient and productive they become [Marr2018].

This leads us to Industry 5.0, assisted by technology such as IoT and Big Data to create collaborative working relationships *between* humans and machines [Özdemir&Hekim2018]. The need for this relationship stems from the fact that while smart machines are better at strenuous and precision work than humans, human intelligence is better at assessing and adapting to unexpected situations than the problem-solving capabilities of a smart machine [Jardine2020]. Industry 5.0 is said to combine both the mechanical efficiency of machines with the creative intelligence of humans [DemirEtAl2019]. In this way, the autonomous robots of Industry 4.0 become the collaborative robots of Industry 5.0. This collaboration mirrors the predications of Web 4.0, using the same notion of matured semantic and reasoning technologies that allow for such a relationship between human and machines.

A result of this new harmony between human and machine workers could be an ability to deliver on consumers' increasing demand for personalisation in the products that they buy. Large scale manufacturing of the kind we currently see in Industry 4.0 is good at the high-speed mass production of a standardised product but is less well-equipped when a consumer requires a customisation to this standard [Østergaard2017]. Conversely, while human workers have the creativity and flexibility required for personalisation, the time this process takes might not always be economically feasible for production. It has been predicted that the collaboration between human and machine could help stimulate large scale personalisation of products by combining the manufacturing proficiency of smart machines and the creativity of their human counterparts [Atwell2017].

Importance of digital identity

The scenarios described in the previous sections highlight the increasing importance of digital identity, given the rising volume of Internet-connected devices and machines [Nicolás2021]. An identity is defined as a "set of attributes related to an entity" [ISO/IEC24760].

A digital identity is information used by digital systems to identify a defined subject, be it a person, machine or otherwise. Digital identity is important to establish trust amongst people and machines, be it in the real world or the Metaverse. Much like the human identities we rely on to access apps and devices we use every day, machines require a set of credentials to authenticate and securely connect with other machines, devices, and apps on the Internet. Clearly, everyone and everything needs it: consumers, machines (manufacturing plants, real-life robots and Metaverse avatars), commercial enterprise, and governments [Callan2021]. Trust in technology was not included as part of the foundations of the Internet[21]. Currently, we mainly rely on centralised approaches, such as those based on shared secret[22] or digital credentials through a third-party. There is ongoing research for other approaches which offer increasing degrees of protection against fraud as well as being scalable and flexible [Windley2018].

Recent approaches to digital identity have been based on the concept of self-sovereign

[21] 'On the Internet nobody knows you are a dog' are the words from a 1993 cartoon penned by Peter Steiner [Steiner1993].

[22] Such as a username and a password.

identity (SSI) [López2020], which advocates that an individual should own and control their identity without the need of third party(s) thus allowing people to interact in the digital world with the same freedom and capacity for trust as they do in the offline world [PwC2019]. SSI not only overcomes existing problems but also gives back to the consumer full control of their identity [PwC2019].

Ledger technology[23] provides one approach to implement practical SSI [IFTF2017].

Importance of Interoperability

Interoperability is a foundational principle of the Internet [O'Hara&Hall2021]: "the Internet's secret sauce, and integral to how it was built" [Riley2019] and it is what makes the Internet as we know it today.

Ideally, people should be able to transfer data/content within and between organisations seamlessly, enabled by software and hardware technologies. Application programming interfaces (APIs) are key to achieving seamless data interoperability both within individual organisations and, when enabled by the respective organisations, between organisational boundaries[24]. In addition to the APIs, data structures and standardisation will also be required for interoperability.

Interoperability is impacted by four factors: technology, data, people, and organisations. Interoperability of data transport technology is essential to facilitate basic connectivity and the transport of the 1's and 0's across the Internet[25]. Interoperability of transport technology is achieved through the on-going standardisation effort of the Internet Engineering Task Force (IETF).

Data interoperability is covered below[26].

Data interoperability

In order to create autonomous operations on the Internet, there is a need to establish semantic interoperability: systems must be able to exchange data in such a way that the precise meaning of the data is readily accessible and the data itself can be translated by any system into a form that it understands. In this way, data can be exchanged and interpreted unambiguously between different systems. The de-siloed information enables each system to benefit, while in some cases allowing the systems to work together in a collaborative fashion.

Consider the case of IoT, where different vertically integrated systems exist. Semantic interoperability would allow the collaboration of all systems enabled by the flow of information between them. At present, the Internet protocol has helped in standardisation at the network level, while long-term efforts in standardising IoT data models have led to several established standards being developed[27]. However, these standards are not compatible when integrated at a high level because of differences in semantic models (e.g. vocabularies and concepts). Furthermore, as the complexity and number of subsystems

[23] Distributed databases that secures, validates and processes transactional data (e.g. blockchain).
[24] If data is equivalent to oil, then APIs are the pipelines that transport it.
[25] Interoperability of data transport is essential to data and semantic interoperability.
[26] People and organisation interoperability are outside the scope this report.
[27] For example, Open Mobile Alliance SpecWorks and IPSO Alliance [OMASpecWorks], Open Connectivity Foundation [OpenConnectivityFoundation], oneM2M [oneM2M], and Zigbee [ZigbeeA] Dotdot [ZigbeeB].

increase, the issue of wider interoperability exacerbates [IEC2019], thus making semantic interoperability a key enabler for unlocking the potential of IoT as well as scaling up relevant applications such as smart cities[28].

The industry has recognised the importance of this challenge and formed the One Data Model (OneDM) initiative to create a joint interoperability framework between existing IoT data model standards [Widell&Keränen2020].

OneDM has introduced the development of the Semantic Definition Format , a neutral mechanism that enables the translation between and alignment of the data models from different ecosystems, see Figure 10.

It is also working on a common set of data models for different IoT devices over time[29]. There is an anticipation that this common set could reduce fragmentation in the future and even further lower the bar for interoperability.

Figure 10: The OneDM model. Reproduced with permission from Ericsson [Widell&Keränen2020].

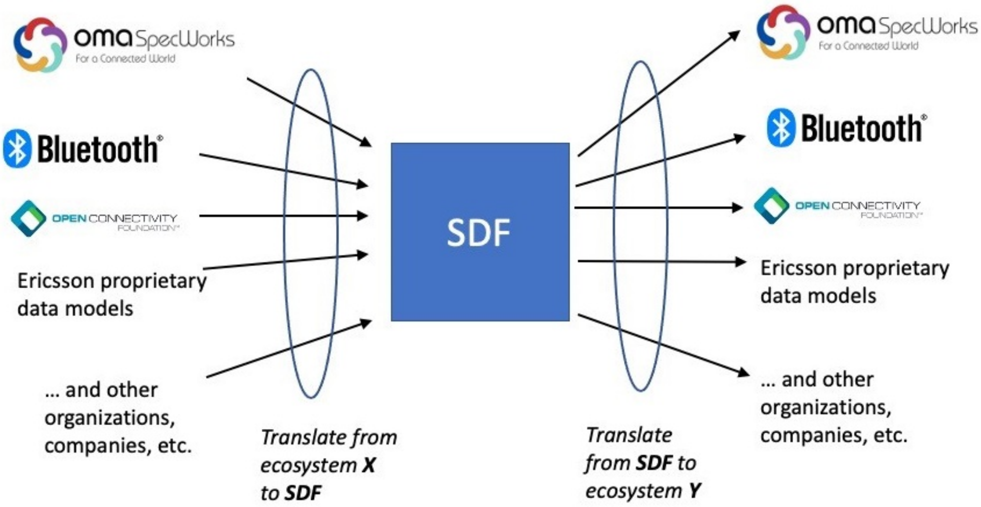

Platform interoperability

Most online social network (OSN) platforms are implemented in a proprietary fashion, keeping their users from seamlessly connecting to and communicating with users of other services. The lack of inter-platform interoperability has led to the current situation of 'isolated islands' where consumers can only communicate within the platform they connect to and are unable to communicate across different platforms. While some platforms offer consumers the option to download their content or data, content/data portability is very difficult in general, making this feature sub optimal.

Concepts such as that implemented by the Social Network Inter-Connect project allow different kinds of OSNs to interoperate using open data formats, common APIs and protocols [SONIC[30], Goendoer&Sharhan2018, DataTransferProject]. These are the key elements that allow the exchange of consumer

[28] Ontologies are intended for semantic interoperability.

[29] Participating organisations such as Open Mobile Alliance SpecWorks, Open Connectivity Foundation, Zigbee Alliance and Bluetooth SIG [Bluetooth].

[30] Not to be confused with the SONiC in chapter 3.

information across platform borders in real-time thus resulting in a heterogeneous network of loosely-coupled OSN platforms, which allow seamless communication between the constituent different platforms [Wong2021]. This way, users can choose the OSN platform they prefer while staying seamlessly connected to their friends using other OSNs.

Towards a ubiquitous Internet

Looking deeper into the combination of the Metaverse, smart cities and Industry 5.0 spheres, suggests that in the future, everything that can be connected could be connected: personal items, objects all around us, our bodies, etc. Further, the pivotal role played by the Internet in this future landscape requires the Internet to be available everywhere. Providing reliable QoS is key to delivering reliable Internet within these three spheres and with good QoE: not only is reliable QoS necessary for the IoS delivery within the realm of the Metaverse, it is also an important factor in supporting the delivery of IoT within smart cities and Industry 5.0.

In general, an application QoS requirement is expressed in terms of two metrics: throughput[31] and latency, which quantify the needed minimum data rate (in bits per second) and the maximum tolerable latency (one-way delay) between the two communicating ends. These are in addition to reliability, trust, energy consumption, amongst other equally important requirements such as the importance of the reliability, availability and consistency of services.

Different applications require different levels of throughput and response for them to work correctly. Indeed, some applications, such as high-quality video applications, require high throughput while others, such as gaming applications, place their demand on low latency[32]. Those that require both high throughput and low latency at the same time are the toughest to deliver.

We have studied the QoS requirements for both the IoS and IoT which might underpin the Web 4.0 and beyond. Figure 11 describes the requirement number of applications in terms of their expected throughput and latency [SofiaEtAl2021, PapadopoulosEtAl2021, Ofcom2021, 5GACIA2019].

The requirements suggest that tomorrow's networks need to be able to deliver end-to-end throughputs in excess of 1Gbps and latencies as low as 100μsec[33] in order to deliver the immersive applications within the context of the IoS and IoT. Clearly, the future need is for high bandwidth and low latency.

Such delivery will require a combination of technologies, including fibre-networks, cellular-networks, Wi-Fi access points and satellites [ForgeEtAl2018] as well as a mixture of public and private networks and cloud-edge computing.

[31] Throughput is the amount of data delivered per second.

[32] Response is the amount of time taken for an application to complete its information exchange with a server.

[33] μsec is read as microsecond which is 1 millionth of a second.

Figure 11: Indicative quality of service requirements for future applications (IoS and IoT). Source: Ofcom.

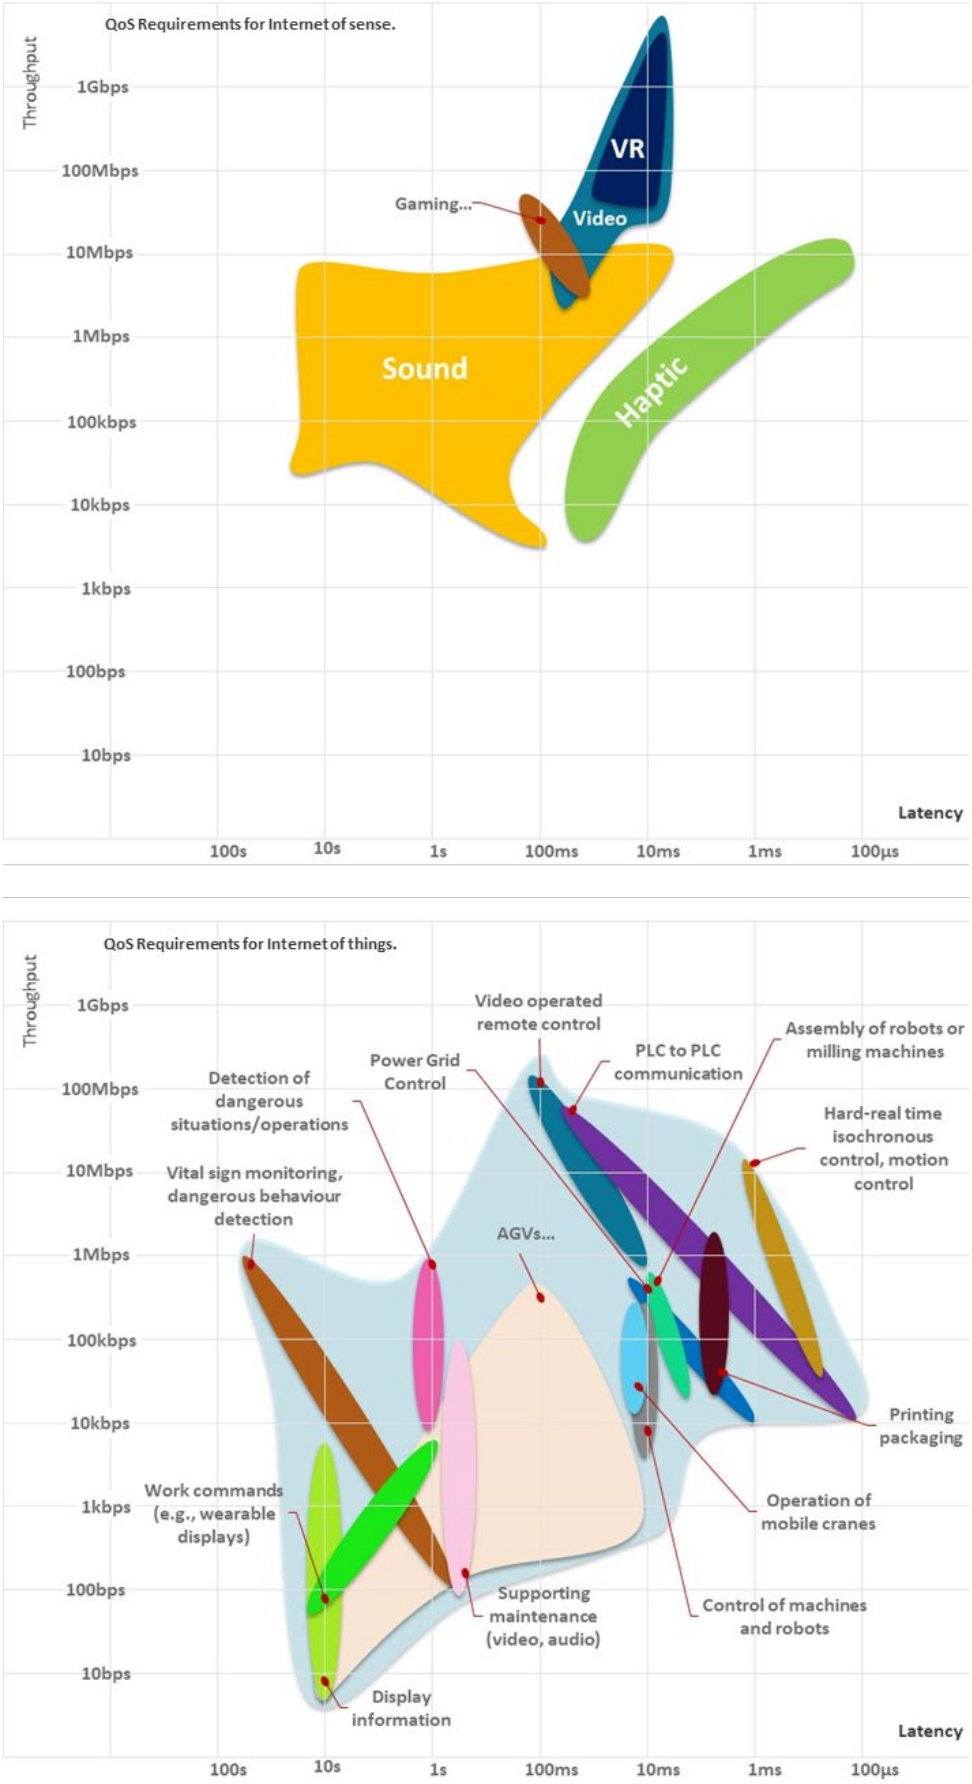

2: Evolution in Internet and Web technology

This chapter addresses technological advances in the Internet and web. We first review recent developments in the area of semantic web and in particular the area of information sharing with trust and privacy. Next, we report on two developments that also address consistent QoS coupled with privacy and security. This is followed by a look at technology development that can improve consumer experience over the web.

We next look at the emerging technologies in Internet suite of protocols that underpins the Internet and hence the web[34]. We first report on the transport layer with recent developments in the Quick UDP Internet Connection protocol (QUIC) that guarantees end-to-end delivery, network congestion control and the key application HTTP/3. This is then followed by the network layer, where we look at Internet Protocol version 6 (IPv6), routing over wireless networks, and interoperable alternatives to IP protocol. We also look at recent advances in routing which is leading to increasing levels of resilience.

Context

Consumers' Internet use in the last decade has highlighted the notion that greater value from the Internet lies in its ability to provide for people sharing of information. One reason for this is the significant role played by the web. While the traditional role of the Internet as the workhorse that delivers traffic continues to be key, the web's ability to provide access to information; and data storage and processing is equally as important.

Consumer and business expectation are for consistent QoE in addition to reliability, privacy, trust, security, and online safety.

Information sharing with trust and privacy

More and more data are being generated in the world. Figure 12 shows that global Internet traffic has been rising [Kamiya2020].

Figure 12: Annual global growth in global Internet traffic. Source: [Kamiya2020][35].

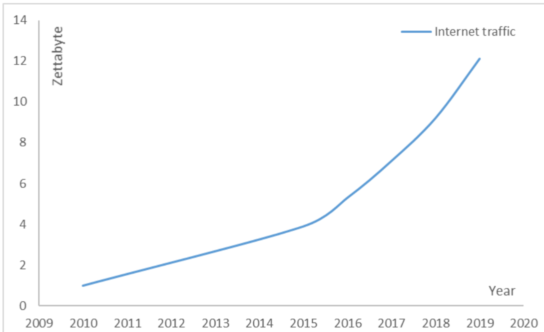

The notion of vendor-owned 'isolated islands' where consumers and/or enterprise can only communicate within individual platforms was discussed in the context of interoperability.

[34] Refer to footnote (6).

[35] Based on IEA data from the IEA (2020), Data Centres and Data Transmission Networks, IEA, Paris, https://www.iea.org/reports/data-centres-and-data-transmission-networks. All rights reserved: as modified by Ofcom.

Such platforms now store vast amounts of siloed data, which has raised questions about privacy, ownership, and innovation. This may also make it hard for the users to access their own data and have control over it.

The semantic web is key for enabling machines and people to work in cooperation[36]. It is an extension of the current web in which information is given well-defined meaning [W3CSemanticWeb].

There is ongoing progress in the area of semantic web. Two examples of such effort are Social Linked Data and Automated Intelligent Regulation platforms which enable trusted sharing of information with privacy for individuals and enterprise, respectively. We discuss these two pieces of work below.

Social Linked Data (Solid): infrastructure for personal data

Social Linked Data (Solid) is Sir Tim Berners-Lee's solution to the challenge of isolated islands, which revolves around the notion of personal data sovereignty [SolidMIT2017, SolidProject, Inrupt]. It is a decentralised platform for social applications on the web, in which users' data is managed separately from the applications that create it and those that consume it. A key technical ingredient that underpins the Solid solution is the concept of the Personal Online Datastore (POD), which would form a special-purpose back-end storage. Effectively, this means each person would own a pod over which they alone have control, and which contains their own data — websites visited, credit card purchases, workout routines, music streamed, etc[37]. Access to a person's data could be gained, with permission from its owner. Such permission may be granted through a secure link for a specific task like processing a loan application or delivering a personalised advert. A company could link to and use personal information selectively, but not store it [Harris2021, Lohr2021]. Consumers are equipped with a range of web-accessible apps that rely on access to pod-stored data with consumer permission. Pods and apps may be provided by a market of independent vendors, thus enabling simple switching between pods and app providers.

Personal data sovereignty is achieved by the fact that a person's data can be linked to but not stored, meaning any permitted application is unable to push such data to a centralised data warehouse.

The Solid solution is based on World Wide Web Consortium recommendations, which are open, in interoperable standards, allowing developers to create applications which can read or write to pods, or control access securely [W3C2020]. Such applications can range over all the data under the user's control wherever it is stored on the web.

Automated Intelligent Regulation: infrastructure for enterprise data

The Automated Intelligent Regulation (AIR) Platform is a data infrastructure that sits on top of the Internet and supports safe sharing of private data [Treleaven2020]. It relies on concepts and technologies such as data sovereignty (privacy, trust, etc), data management (collaboration, security, monetisation, etc), and data technologies (universal digital object identifiers, and

[36] See chapter 1 for a discussion on web evolution.
[37] A person can have as many pods as he/she wants or needs. Also, they may outsource POD management to a third party, while they retain ultimate authority over their PODs – in this case, they may not be involved in day-to-day management of their PODs.

machine learning, data models, blockchains, etc).

The AIR Platform combines multiple next generation technologies to deliver this new infrastructure. Its foundation is distributed by the ledger technology[38] in combination with data security. It also relies on the data layer[39] which conforms to common data standards, often for a specific industrial sector. Universal digital object identifier technology is used to provide unique, persistent, and resolvable addresses for the data objects (document, policy holder, company, etc). Federated learning allows machine learning algorithms to be trained across distributed and isolated data sets, with the potential for privacy preservation [DW2020].

The AIR Platform could be applied to smart cities, Industry 5.0, healthcare, financial services and amongst other areas discussed in chapter 1.

Enabling consistent security, privacy and QoE on the move

As we roam around, our mobile devices connect to multiple networks to maintain our online connectivity. While this gives access to a variety of valuable services, these network connections are increasingly coming under scrutiny from trust, security, performance, and availability perspectives.

Personal Virtual Networks (PVNs) have been proposed to address these challenges [Choffnes2016]. A PVN is a software-defined network which is enabled and owned by one consumer, who own the security and privacy controls of their PVN as well as other aspects such as their media server, personal information repositories, etc. Being software-defined means that one's PVN can execute on a home router, inside an access network, or in any cloud environment.

Coupled with a ubiquitous cloud-edge infrastructure (see chapter 3), it means a PVN can be made to 'follow' its owner where they go. This will not only maintain the same security and privacy specifications, but also maintain consistent owner QoE [Rashid2020].

Connectivity between the PVN and the consumer devices is established using a secure and flexible interface (e.g. a Virtual Private Network) and the PVN itself is secured from external threats. Consumer trust of the PVN is derived from the fact that its configuration and control are under their own management and can be verified.

Messaging layer security for more efficient and secure distribution of multi-media

Modern-day messaging applications are expected to provide end-to-end security, guaranteeing secrecy of message contents to anyone but the intended recipients. Cryptographic protocols based on pairwise security protocols[40] where two-party secret key exchange have been used over the years for this purpose [Green2018]. However, extending this pairwise security protocols to group communications leads to significant inefficiencies [BeurdoucheEtAl2019].

Message Layer Security (MLS) addresses the cryptographic inefficiencies in large groups (thousands even) in messaging, audio, and video calls; and has a direct impact on improving quality [BeurdoucheEtAl2019, IETF/MLS]. With pairwise key-exchange, each

[38] See footnote (23).
[39] See footnote (6).

[40] Security protocols that require no third-party for exchange of security related information.

communicating end in groups of N using pairwise key-exchange case has to store (N-1) keys, and encrypt and forward the same message (N-1) times. In contrast using MLS, each communicating end would need to store one key; and encrypt and forward the message once to a distribution server, see Figure 13.

Figure 13: MLS reduces the number of times a message needs to be encrypted (e.g. from 4 to 1, as shown). Source: Ofcom.

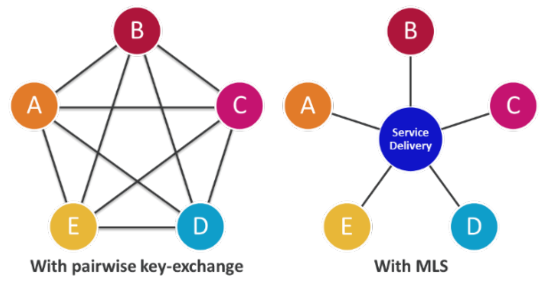

Thus, MLS decreases size and processing time of messages and the resultant transmission traffic which leads to better user QoE and energy usage[41].

The creation of the new open Internet Engineering Task Force (IETF) standard enables interoperability within the sector, where different vendor multi-media applications can be inter-connected in a secure, federated fashion.

WebAssembly for better web experience

In the context of web applications, the client/server paradigm[42] refers to any application which is wholly or partly hosted within a web browser (client). The browser makes requests of the web server for information. On the server side, web server software processes the requests and returns the requested information to the client, which in turn displays the information in the browser on the screen. In its response, the server uses a standardised language for creating web pages known as Hypertext Markup Language (HTML) [W3Schools][43]. Programs (scripts) may be embedded within the HTML. Indeed, a website is a group of client-side and server-side programs (scripts), and HTML.

A server-side script is invoked on the web server when the client browser requests information. A client-side script is a program that is processed within the client browser. These scripts are programs which are downloaded and run by the browser. JavaScript is an important and widely used client-side scripting language. Over the years, scripting languages have been optimised in order to improve the consumer experience when web browsing. On the browser side, JavaScript execution engines have been streamlined and has received a great deal of attention by browser vendors in order to achieve faster execution.

Since being introduced to enable the running of large complex applications on the web, client-side WebAssembly has become an industry standard [W3C2015]. WebAssembly is an execution platform created with efficiency, safety, and portability in mind. Its main goal is

[41] For example, consider 100 thousand members and message size of 1 kbit (equivalent to 128 English characters). Pairwise based solution would require 100k operations and payload of 100 000 × 1 kbit = 100Mbit. In contrast with MLS, the same message will require 17 operations and payload of 17 × 1kbit = 17kbit.

[42] Computers that provide services on the Internet are called servers. A server accepts connections from other devices on the Internet to send information to, receive information from or carry out processing for them. A client is the computer, or user who makes the connection request to the server.

[43] HTML may contain image files (like GIF and JPEG), and object files (like sound and video).

to enable high performance applications on the web. Its key goals are to get near-native performance without using an add-on and executes code within the browser's in an isolated fashion to achieve both better performance and safety. Thus, WebAssembly turns web browsers into a high-performance virtual machine for the open web. It makes it possible to support in-browser, graphics-heavy games, port scientific simulations, and other compute-intensive applications.

WebAssembly is being added to most browsers, see Figure 14. Besides JavaScript, it enables browsers to run programs written in computer languages such as Rust and C/C++, with additional languages support to come, thus allowing the web to grow beyond JavaScript.

Figure 14: Overview of WebAssembly. Source: Ofcom.

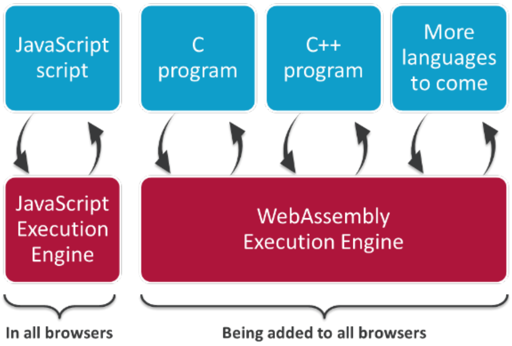

More recently, developers have realised that WebAssembly has the potential to transform software development beyond the browser by overcoming the security and stability challenges of relying on untrusted software from third parties[44]. This has triggered industrial effort to evolve WebAssembly for server-side, see chapter 3 for further discussion.

The adoption of WebAssembly on the client and server sides brings improvement in web QoE through improvement both in the web response and an increase in the richness of applications.

Developments in the transport

To deliver the services discussed in chapter 1, the Internet community has been exploring several directions for further evolving the transport protocols. These protocols are pivotal in the delivery of key services such as the web and many other applications.

QUIC

After five years of collaborative work, a new transport protocol called QUIC has recently been standardised by the IETF [IETF-RFC9000, Katz2021]. QUIC reduces connection establishment latency by integrating connection setup and security handshake [Kumar2020]. It also reduces overall server response. It has been designed to use multiple streams and run over the connectionless User Datagram Protocol (UDP). QUIC traffic is encrypted including transport headers, see Figure 15. Deployment of QUIC has already started and reports indicate that the IETF QUIC significantly outperforms HTTP over Transport Layer Security (TLS) 1.3 over Transmission Control Protocol (TCP) [SchinaziEtAl2020]. Yet, the implementation is expected to take some time to mature and attract the support required to spread widely.

[44] Tools like containers can provide some degree of isolation, they also add an overhead.

Figure 15: QUIC and its relation to other IP layers. Source: Ofcom.

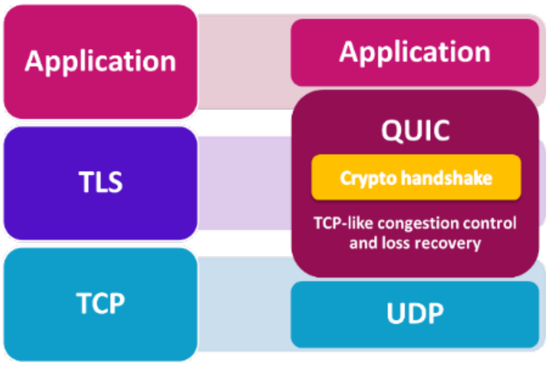

Another proposed development for QUIC is multipath extension for QUIC (MP-QUIC) which is a current IETF work in progress [IETF-MPQUIC]. One of its potential benefits is seamless use of Wi-Fi and mobile network for smartphones.

Congestion control

One fundamental element of TCP and/or QUIC is congestion control, which manages traffic flow when the network becomes congested. Over the years, improvements have been made through the generation of different versions of congestion control algorithms such as TCP Tahoe, TCP Reno, TCP Vegas and TCP CUBIC [Duarte2008], which is aimed at higher throughput and embedded security.

A TCP connection is affected by the slowest link in each direction at any given moment. Usually the slowest link is affected by congestion, and is where persistent queues form. Algorithms like CUBIC rely on packet loss as an indicator for congestion. In contrast the Bottleneck Bandwidth and Round-trip propagation time (BBR) algorithm periodically estimates the round-trip time and the available bandwidth of the flow.

Relative to loss-based congestion control algorithms such as Reno or CUBIC [IETF-RFC5681, IETF-RFC8312], BBR offers substantially higher throughput for bottlenecks with shallow buffers[45] or random losses, and substantially lower queueing delays for bottlenecks with deep buffers[46], thus avoiding 'buffer bloat'[47]. This algorithm can be implemented in any transport protocol that supports packet-delivery acknowledgment[48]. BBR has been gaining widespread attention because of its ability to operate without creating packet loss or filling queues. This algorithm can be implemented in both TCP and QUIC and will continue to be optimised and evolved in the future as its adoption is spread more widely in data centres and, eventually, into Internet browsers [CardwellEtAl2021].

Low Latency, Low Loss, Scalable Throughput (L4S) is a new technology currently under discussion at the IETF [IETF-L4S]. It is aimed at allowing the use of a new class of congestion controls that ensures low latency and high throughput for Internet applications by reducing the TCP packets' queuing latency while maintaining high bandwidth. The new class of L4S congestion controls can coexist with the other congestion controls in a shared network without impacting the performance of the other non L4S traffic.

L4S could provide better QoE underload for applications such as gaming and video conferencing. It also allows applications such as gaming, interactive video and, Augmented

[45] Almost every network device contains a 'queue' or alternatively known as a buffer, which is a piece of memory where network packets are stored as they await processing time. A shallow buffer refers to a queue that store a small number of packets.
[46] A deep buffer refers to a queue that can store a large number of packets.
[47] Buffer bloat describes the situation when congestion occur in the network causing packets to be queued for a long period of time and thus leading to large delays and ultimately poor QoE.
[48] Open source implementations are available for TCP and QUIC [IETF-RFC793, IETF-RFC9000].

Reality and Virtual Reality to migrate into the cloud.

HTTP/3

Since around 2015, the HTTP/2 protocol has been deployed widely across the Internet and the web and addresses many of HTTP/1 shortcomings. HTTP/2 makes use of many logical streams that are sent over by the same physical TCP connection, making data transfer faster [IETF-RFC7540].

HTTP/3 is a new reliable and secure transport protocol which will make use of QUIC [IETF-http3]. It aims at addressing some of the known shortcomings of HTTP/2 over TCP and TLS. It is expected to provide better consumer QoE. For example, it will benefit from QUIC's zero round-trip times (0-RTT)[49] handshakes which are much faster than the combination of TCP and TLS [IETF-RFC8446]. Also, HTTP/3 only exists in secure, encrypted version.

Developments in the network

The most important protocol in the Internet layer is the IP. It is the workhorse that carries traffic across the globe and is the standard for routing packets across interconnected networks.

IPv6: more adoption and more enhancements

The global penetration of the Internet has led to the exhaustion of IP version 4 (IPv4) addresses. By 2018, no more unallocated, globally unique numbers could be issued as the original Internet's 32-bit address space had all been allocated. IPv6 is the latest version of the IP protocol and is designed to overcome IPv4 issues, such as limited address space [IETF-RFC4291].

The signs are heading in the direction of IPv6 as more IPv6 only applications emerge [IETF-IPv6-deployment]. This may be stimulated by the proliferation of IoT devices. The adoption of IPv6 by consumer IoT initiatives such as 'Connected home over IP' may very well help to speed up mass adoption [ZigbeeC, BW2021]. Similarly, industrial IoT is already heavily dependent on IPv6. It may take a while until we see full IPv6 adoption, with IPv4 and IPv6 likely to continue running in parallel[50] in the near future.

IPv6 continues to receive enhancements and optimisation by the IETF. One key driver is the evolution towards application centric networks and platforms, which are raising operators' requirements for more flexible, yet scalable, and simple to operate network architectures, packet forwarding and traffic engineering practices. Segment routing over IPv6 (SRv6) is a solution for IP backbones and datacentres; and can extend beyond as well [FilsfilsEtAl2015]. It facilitates dynamic real-time configuration (programming) of intermediary active elements (network switches and routers) based on information found inside the packets which describe the end-to-end network QoS requirements.

Reliable and Available Wireless (RAW)

Deterministic Networking (DetNet) aims to enable the traffic of applications with critical timing and reliability issues to be delivered over IP, and to support both the new applications and existing packet network applications over the same physical network. To this end, it reserves network resources in

[49] Reduction in the number of TLS acknowledgement needed once an initial session is set up and this dramatically speeds up resumed connections to a webpage and hence leads to a better QoE.

[50] This is known as dual-stacking.

order to carry unicast or multicast data flows with bounded delay and jitter, ultra-low packet loss, and ordered delivery; and on a per-flow basis [IETF-DetNet].

The IETF is developing Reliable and Available Wireless (RAW) technology which is an approach for deterministic networking over wireless paths. Because of the non-reliability of wireless, RAW addresses the additional issues of less consistent transmission links [IETF-RAW].

RAW is being considered for carrying non-IP traffic such as High-Definition Multimedia Interface, Controller Area Network, etc as well as industrial automation and IoT; professional audio and video; gaming; aeronautical data communications; and edge robotics.

RAW uses link state information to compute physically link and node-disjoint paths in the network. This information is used for route computation through the network and the identification of routes that meet the required end-to-end QoS.

Actual traffic delivery is achieved using a distributed routing paradigm that overcomes the potential issue of slow centralised control. Thus, the RAW Path Selection Engine performs rapid local adjustments (sub-second) of the forwarding tables and uses advanced IPv6 technologies such as SRv6 to actually steer traffic across redundant network paths in order to achieve high reliability and meet QoS requirements.

Semantic addressing and routing

With semantic addressing, an IP address can have a different meaning beyond its network location [King&Farrel2021a]. Examples of semantic addressing include multicast addressing where the assigned prefix can be associate with the content type (TV, radio, etc), semantics for mobile users so that they can move without facing service disruption and semantics to include geographic location information within the address [King&Farrel2021b].

Semantic routing is a way of routing packets based on semantic addressing and on additional information carried within the packets (other than traditional least cost path routing) [King&Farrel2021a].

Introduction of new address semantics could help overcome limitations of the current technologies e.g. the use of space geo-location address to support highly dynamic network topologies like the low earth orbit satellites. New IP address semantics can also be used to provide better QoS and efficiency during the transportation of packets [Galis&Lou2021].

Several semantic routing solutions have been proposed [King&Farrel2021b]. Some of the solutions are proposed for deployment in 'isolated islands' which may introduce interoperability challenges. Other challenges of semantic routing in IP networks include risks of privacy, stability and scaling problems for the routing protocols, and compatibility with existing networks [King&Farrel2021a].

Early discussions have started within Internet community to look at future approach to semantic addressing and routing in terms of use cases, requirements for future networks and the associated challenges [Galis&Lou2021].

Interoperable alternatives to IP

The original IP suite of protocols were not designed for mobility, security and the provisioning of QoS. Over the years, the Internet community has made a great deal of effort to cater for such requirements in both the network and transport layers by introducing protocols such as IPv6 and QUIC.

The format of the IP packet header is fixed, and the header carries all the packet forwarding information which requires hardware support to process it. Introduction of new packet formats may take an exceedingly long time as new hardware will be needed. To overcome some of the inefficiencies of IP and support new use cases such as those discussed in chapter 1, the Internet community have been studying more efficient alternatives to IP protocols. An example of this is discussed further below.

Non-IP Networking: Recently, the European Telecommunication Standards Institute (ETSI) launched a new group on Non-IP Networking (NIN) which intends to develop new protocols that are more suitable for the future and with improved security, efficiency, and performance [ETSI-NIN, Antipolis2020]. In general, NIN protocol aims to:

1. Make it easy to introduce new addressing scheme because of the separation of the forwarding and control planes.
2. Provide better QoS for delay sensitive applications and security.
3. Provide better efficiency in terms of spectrum and processing power.
4. Support multicasting for unlimited number of users.
5. Provide better web performance.

One candidate of NIN technology is FlexiLink which supports both IP and non-IP traffic, guarantees timing, and is efficient in terms of overhead [ETSI2021].

FlexiLink sends packet and routing information separately, and this reduces packet size considerably compared to the IP packet size, which makes it easier and more efficient for the network switches to forward packets. FlexiLink further simplifies packet forwarding by adding a label to the packet header that is used as an index in the routing table. In FlexiLink a stream of packets is considered as a flow and control messages are sent per flow rather than per packet. Again, this is claimed to increase efficiency.

FlexiLink may find application in private networks where very tight control on end-to-end latency is critical, e.g. robots running around the factory.

Towards resilient Internet routing

Border Gateway Protocol (BGP) is one of the key protocols for the Internet, as it provides the means to connect different service providers' networks, or Autonomous Systems (ASs)[51]. Like other aspects of networking, this protocol is also at risk of attacks which can harm individual users and/or disrupt network operations.

Resource Public Key Infrastructure (RPKI) is a public key infrastructure framework designed to secure the Internet's routing infrastructure, specifically the BGP [Durand2020]. Key to RPKI is the attestation of a BGP route announcement (Route Origin Authorisation), which provides greater assurance that the origin AS number is authorised to announce the prefix(es); and the attestation can be verified cryptographically using RPKI. Thus, legitimate holders of number resources are able to control the operation of Internet routing protocols to minimise the impact of malicious activities, such as origin prefix hijacking[52].

[51] Each Asynchronous System is given its own unique identifier (number).

[52] Traffic will not reach its intended destination when the BGP route to that destination is wrongly announced. This could happen intentionally by a malicious unauthorised party.

While adoption of RPKI throughout the Internet is gaining momentum, it still cannot provide end-to-end BGP path security and resilience. Additional efforts are being made by the Internet community to secure the end-to-end BGP path routing. An example of a recent effort is discussed further below.

Scalability, Control, and Isolation on Next-Generation Networks (SCION): this is a clean-slate Internet architecture designed to provide route control, failure isolation, and explicit trust information for end-to-end communication. It is said to achieve end-to-end resilience of the Internet routing system [Hitz2021, De-Ruiter&Schutijser2020].

The SCION network consists of interconnected isolation domains (ISDs), each of which consists of an interconnected set of ASs. Each isolation domain defines its own PKI which could be used to secure routing within the ISD so that the effect of routing failures can be limited to the affected ISD. Each AS also needs a certificate to be a member of an ISD. This enables access control and policy enforcement for ISDs.

Application-aware networking and Network-aware applications

Traffic engineering is primarily intended to allow differentiated treatment to be applied to certain types of traffic to meet QoS requirements [IETF-RFC3272].

Application-aware networking allows application characteristic information such as application identity and its QoS requirements to be carried in the packet encapsulation. This allows networks to adjust in response to the application characteristic information to meet the application QoS requirements [LiEtAl2020].

To achieve the required consumer QoE, the network must first provide fine-granularity and perhaps application-level guarantees. Then, the application information is used by the network to steer the traffic through the appropriate paths, push the traffic through certain queues or even allocate dedicated network resource to meet the application requirements.

There is also work on increasing application awareness of networks. The goal of application-layer traffic optimisation has been to provide guidance to applications, for example to select one or several hosts from a set of candidates that are able to provide a desired resource [IETF-RFC7285]. The concept is now being extended to include making network performance metrics available to applications on request [WuEtAl2021].

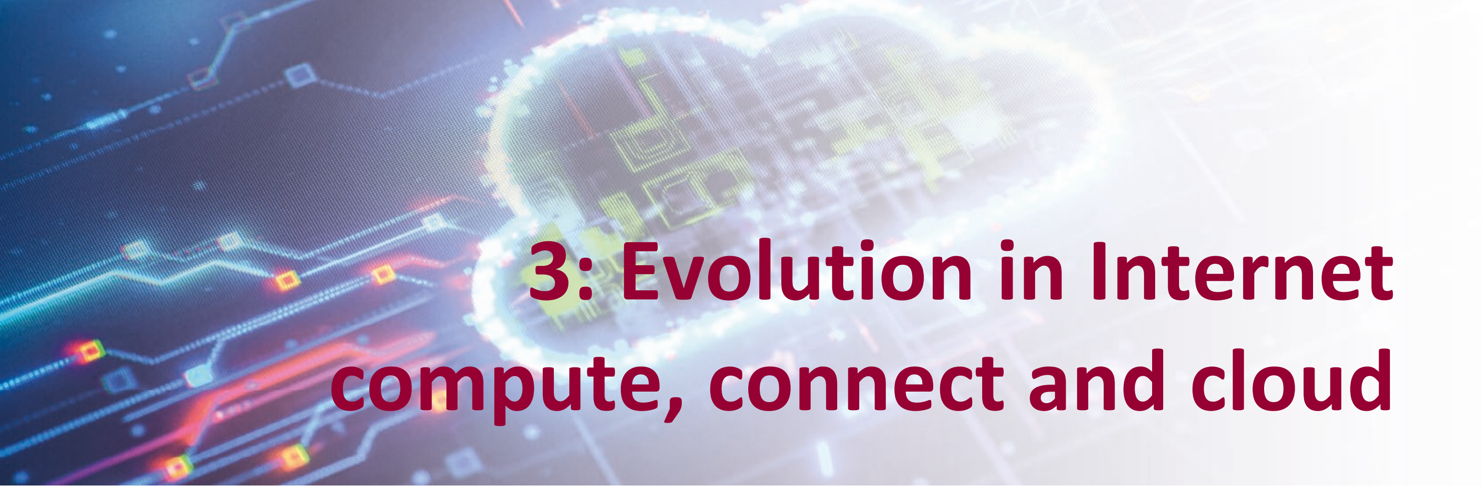

3: Evolution in Internet compute, connect and cloud

Chapter 1 indicated that the proliferation of more IoT devices and rich applications would generate an increasing amount of information and traffic. At the same time, consumers and businesses expect better and more consistent quality.

This chapter looks at technology developments which will help process large amounts of data at a higher speed and help deliver it to consumers and businesses.

Context

The continued consumer and business expectations call for increasingly high-performing, faster computation and communication technology in order to instantly access information; and/or perform data processing and return the results. This high-performance is split into broadly three areas:

1. Compute technology that carries out the computation;
2. Connect technology which moves data between two communicating ends; and
3. Cloud technology, which brings processing and storage into one framework to cater for consumer and business needs.

Compute technology

We look at two aspects of computing technology: semiconductors, and processors and computing architectures, which are the basis for data computation at the computer system level.

Semiconductors

The foundations of processing and storage, and networking are built on semiconductors in the form of transistors. Traditionally, the number of transistors is coupled with the number of calculations a chip can perform.

Four years ago, IBM created an architecture designed to produce a chip with transistors as small as 5nm[53]. Building on that work, IBM reported in 2021 the production of the world's first 2nm transistors. Each 2nm-based chip is the size of a fingernail and contains 50billion 2nm transistors.

The 2nm technology is designed to improve calculation speed or energy efficiency. Where calculation speed is the concern, the 2nm design is projected to achieve 45% higher performance than today's 7nm chips. If energy efficiency is the goal, then the 2nm technology is expected to use 75% less energy than 5nm chips [Gartenberg2021, Patrizio2021a]. It has been indicated that 2nm processor technology

[53] A human hair is approximately 80,000-100,000 nanometres wide and a nanometre is one billionth of a meter (10^{-9}).

could be rolling out around 2024 [Johnson2021].

Advances like this help secure the continuity of Moore's law[54]. While Moore's law has principally maintained up to now, innovation in this area has been slowing as the size of transistors reaches the near-atomic level. In 2017, it was reported that time might be running out for Moore's law and predicted that advancements in chip technologies would soon begin to slow [Loughran2017, IET2021].

Processor and computing architecture

The von Neumann processor architecture has been the foundation of processor hardware design. Today's multi-core[55] computing systems are still based on this highly successful architecture. The von Neumann architecture requires data to be moved back and forth between the physically separated processor and memory, thus introducing memory latency. On occasions and for specific applications, a high-speed processor could find itself more in time idle waiting for data to arrive from memory. Various remedies have been implemented (e.g. caching, multi-threading, etc) with varying degrees of success [NN2020].

Another way to speed up the computation process is by adopting alternative processor architectures to the von Neumann [Ofcom2021]. Whilst the general-purpose utility of von Neuman means it will likely remain in use in the foreseeable future, more specialised non-von Neumann computation platforms are being developed as co-processing complements that achieve higher computation speed for certain types applications.

Instruction set architecture (ISA): The fifth generation Reduced Instruction Set Computer (RISC-V) is a license-free, open, and extensible processor architecture with an inherit flexibility, scalability, extensibility, and modularity [Hennessy&Patterson2019, RISC-V].

A key feature of RISC-V is its simplicity[56] and small instruction set, which reduce the effort needed at the designing and hardware verification stages. It can be extended through optional extensions[57] and both 32- and 64-bit versions are possible as well as multicore implementations.

RISC-V processors can be made to act as a general-purpose processor as well as a custom processor dedicated to one application. Customisation for specific applications is achievable by reserving a vast instruction space for acceleration engines. In this way, the customisation offers orders of magnitude improvements in performance, cost and energy when compared to general purpose processors.

RISC-V has been gaining popularity in industry because of the following [Marena2018]:

- The free, open-source hardware and software (permissive license) provides an alternative to closed, costly ISAs.

[54] An observation first made by Intel founder Gordon Moore in 1965 that the number of transistors in a dense integrated circuit would double about every two years.

[55] Including parallel Central Processing Units (CPUs), Graphics Processing Unit, Single Instruction Multiple Data - Very Long Instruction Word (SIMD-VLIWs) and, vector processors.

[56] It has fewer instructions: 50 in the base and the remaining standard extensions add 53 and 34 instructions, totalling 137.

[57] Such as integer multiply/divide, atomic memory operations, single/double-precision floating-point and compressed instructions.

- The availability of freely downloadable cores.
- Implementations can be equivalent in efficiency to alternatives.
- Several Linux distributions are supported.

RISC-V engines have been finding wide-ranging applications in simulation, image processing, deep learning, bioinformatics, digital signal processing, and cryptography.

Being open allows RISC-V ISA evolution to occur in public, with collaboration from hardware and software experts. Similarly, and as an open framework, RISC-V acts a platform for innovation which has formed the basis for collaboration between companies, universities and even individuals to develop special-purpose processors for niche applications or provide alternatives to the traditional x86 architecture.

Recently, the European Processor Initiative (EPI) announced that it has designed its first processor based on RISC-V technology aimed at High Performance Computing platforms, and it is currently having the device fabricated. The aim is to achieve high throughput and low power levels; and support peripherals at rates above 200 Gbps [HPCwire2021].

Memory-centric computing: This approach advocates the use of a large number of memory and processing units combined together so that the computational operation is performed in the same location where data is stored, thus eliminating most of the data transfers between memory and Central Processing Units (CPU). It has been reported that memory-centric computing can speed up processing by 200 times for some applications and with greater energy efficiency [Shepard2020].

The memristor is a potential enabling technology for memory-centric computing [HamdiouiEtAl2015]. It is a semiconductor that joins a capacitor, resistor, and inductor to make a basic building block whose resistance varies as a function of current and flux [Ahmed&Taha2021, Shahsavari2013]. As a basic building block, it can be used to construct logic gates and memory [Maan&James2016].

Memory-centric processing can be used in a variety of tasks ranging from high-precision scientific computing to imprecise stochastic computing, and deep learning in neural networks [NN2020].

Neuromorphic computing: One of the most prominent versions of a memory-centric architecture is neuromorphic computing which is inspired by the workings of the human body, where neurons carry messages to and from the brain [Ofcom2021]. Neuromorphic systems can be either digital or analogue [Soni2018, Intel].

For example, if a person accidently touches hot metal with their finger, receptors in the finger send a message to the brain via a number of neurons. On arrival, the brain registers the pain and sends a message back via a number of neurons to the finger muscle to retract. Many concurrent inputs (spatial) or an input that builds up over time (temporal) can trigger such action. One key aspect is the parallelism involved in the sense that many receptors are working with the brain at the same time both in the spatial and temporal dimensions. Neuromorphic computing mimics the above, where spiking neural networks convey information in both the same temporal and spatial way and so produce more than one of two outputs.

The neurons are the basic computing unit in a human brain, where a vast number of neurons are inter-connected by synapses[58]. Thus, it has been necessary to develop nanoscale, low power, synapse-like devices [JoEtAl2010]. Researchers have been using memristors (discussed earlier) for the fabrication of electronic synapses for neuromorphic computing that mimics some of the aspects of learning and computation in human brains. Interest in memristor devices has also increased because these devices emulate the memory and learning properties of biological synapses.

Neuromorphic computing is evolving and a neuromorphic supercomputer has been created [ManchesterU2018]. It is reported to be capable of completing more than 200 trillion actions per second. It has been used to simulate high-level real-time processing in a range of isolated brain networks of 80,000 neuron model of a segment of the cortex. It has also been used as part of a robot system to interpret real-time visual information and navigate towards certain objects. Generally, application of neuromorphic computing is within the area of AI because of its suitability to deal with probabilistic computing with noisy and uncertain data.

Optical computing: An all optical computing platform manipulates information with light alone because photons, whether generated by lasers or light-emitting diodes (LEDs), can be used to encode data in a similar fashion to conventional transistors in traditional electron-based semiconductors. Computation is normally achieved by adding two light waves coming from two different sources and then projecting the result onto '0' or '1' states. Optical logic gates are possible to construct from multiple beams: one control and one or more to be computed, resulting in a logical output. Optical transistors are slowly being designed and implemented which will eventually allow for an entirely optical computer to be built [EA-Light2021].

Photons are massless and hence, they travel faster than electrons. Also, photon-based transistors are smaller in size than traditional transistors. These properties mean all-optical computers could in principle be capable of operating at a much higher speed than traditional computers.

Quantum computing and Internet: Quantum computers have been the subject of research and development in recent years both within academia and industry because of the unique properties that make them powerful. Quantum computers leverage quantum mechanical phenomena such as superposition and entanglement in order to identify, process and transport data.

Current computers manipulate individual bits, which store and process information as binary ‘0’ and ‘1’. Quantum computers manipulate information by relying on quantum bits, or qubits [Nayak2021]. For example, where a 2-bit register in a traditional computer can store only one of four binary configurations (00, 01, 10, or 11) at any given time, a quantum computer can access a wider set of states, enabled by the superpositions (i.e. combinations of) each of the four traditional states. This increased state space proves useful for certain types of problems, such that computations with qubits can be performed up

[58] Note that a mouse brain consists of around 100 million neurons and the human brain is 1000 times bigger than that.

to a million times faster when compared to traditional bits.

This approach would permit extremely fast methods for solving combinatorial and optimisation problems. In this way, quantum computers are special purpose machines that can solve certain problems faster than existing computers or which may not be resolved by traditional computing because of the length of processing time it would take. They have the potential to impact many fields such as research in security, chemistry, biology, life science, big data, healthcare and telecommunication [Djordjevic2012][59].

Quantum superposition and entanglement provide the foundation for quantum networks and enable data to exist in multiple quantum states, and , among multiple remote quantum devices or nodes [CaleffiEtAl2018]. While many technical barriers still remain, recent results show progress is being made towards building a future quantum Internet with quantum secured information being sent between metropolitan areas [EA-Quantum2021].

Connect technology

Data centres are key elements in the cloud-edge infrastructure we describe in chapter 4. They have historically been built using off-the-shelf fibre communication technology, where it is employed in three ways:

1. Interconnecting the data centres themselves,
2. Interconnecting the various elements within data centres such as the server racks (inter-rack), and
3. Interconnecting the servers within the rack (intra-rack).

To help deliver growing Internet traffic, higher data rates across the optical-electrical interface are needed at a lower power and cost. Increasing the data rate is achievable with more bits per second and greater number of fibres (lane count) as well as the transmission technology such as single- vs multi-mode fibre and the number of wavelengths. However, the optical-electrical interface itself must evolve to facilitate higher-speed conversion from data-carrying optical signal in the fibre to the equivalent data-carrying electrical signal in the silicon and vice versa, see Figure 16.

Bringing optics and silicon together in the same package can save significant power and support higher data rates. Thus, going forward, the co-integration of electrical and optical parts is considered as the next step and co-packaged optics brings the optics much closer to the silicon [Chopra2021]. However, the integration of high-speed electronics and optics may limit the flexibility and interoperability currently enjoyed by the separation of the two, as such, in addition to the technical challenges to be addressed in integration, standardisation will be also be an important element.

Figure 16: Interfacing the optical and electrical domains. Source: Ofcom.

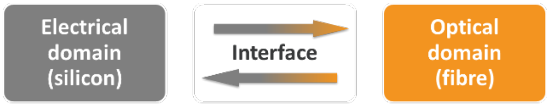

Nevertheless, next generation optical-electrical links are expected to deliver 1.6 Tbps, and a road has been described to deliver speeds of 25Tbps, 50Tbps, 100Tbps and beyond [MinkenbergEtAl2021][60].

[59] Potentially improve techniques of interference and error correction.

[60] Tbit stands for Terabit, which is 1 trillion bits, a Terabit is 1000 Gigabits, or 125 Gigabytes.

Cloud technology

Cloud computing makes computing services (including storage and processing) available to a wide range of users over the Internet[61]. Cloud vendors use virtualisation technology to increase hardware utilisation, allowing workloads from different organisations to be concurrently processed and thus make efficient use of compute resources [ArmbrustEtAl2009].

Cloud technology has been evolving over the years with the aim of increasing efficiency through the introduction of new architectures, different virtualisation types with varying degrees of capabilities or improvement in the actual technology fabric.

The rise of the edge

Edge computing is an umbrella term for a set of architectures and distributed computing techniques that aim to bring cloud compute and storage closer to the consumer than the traditional cloud while maintaining central or distributed control.

Over the years, the Internet has been used to deliver multi-media (video, music, web, etc) data, and enterprise and public information services. This has been achieved through a combination of web technology and the use of Content Distribution Network (CDN)- or cloud-based platforms. Consequently, the Internet has been predominately download-centric. Looking at the applications studied in chapter 1, more traffic will be generated from the end devices leading the Internet to change to a mixture of upload and download centricity. This suggests that the load presented to the core network could increase beyond the capacity of the core. Edge computing can help reduce the amount of traffic travelling across the core and protect it from potential future congestion.

Generally, some of the applications discussed chapter 1 would require stringent QoS that can only be delivered with good QoE by edge computing. Such applications include augmented reality and virtual reality which benefit from lightning-fast response times and low latency communications; connected cars, which needs high-capacity, low-latency, and highly available settings; and other Internet of Things applications such as video monitoring.

As mentioned in chapter 1, different applications require different levels of throughput and response for them to work correctly. Indeed, some applications, such as high-quality video applications, require high throughput while other applications, such gaming applications, place their demand on low latency[62]. Those that require both high throughput and low latency at the same time are the toughest to deliver. Figure 11 describes the requirement number of applications in terms of their expected throughput and latency.

In addition, edge computing could play a role where privacy and security are of concern [OsiaEtAl2020]. Here, the closeness of the edge to the end user, and in some cases in their home or work, offers the opportunity to perform on-premise processing that can ensure data is stripped of any privacy related

[61] The US National Institute of Standards and Technology (NIST) uses the following definition of cloud computing [Mell&Grance2011]: “a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources […] that can be rapidly provisioned and released with minimal management effort or service provider interaction.”

[62] Response is the amount of time taken for an application to complete its information exchange with a server.

aspects before forwarding on towards the core.

Figure 17 summarises the type of future applications in terms of their requirements that calls for edge computing.

Figure 17: Applications types that raise the requirements for edge cloud. Source: Ofcom.

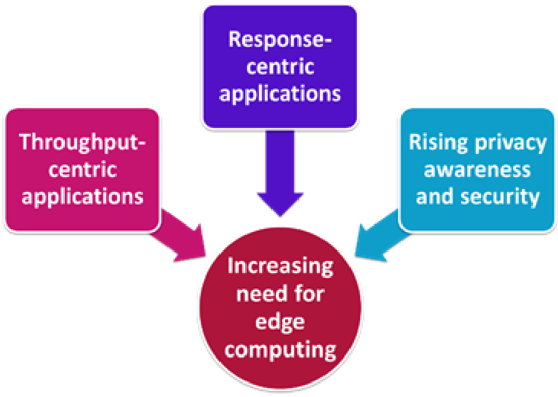

Convergence of cloud and edge

Depending on the scenario, the edge computing application user may be either a human, a business, or a machine (network element, car, robot, sensor, etc). The physical or even logical definition of edge depends on the use case. Example of edge domains include factory floor, residential home, radio access network, operator's core network, operators' data centre network, cloud vendor's distributed data centre network, or the vendor private cloud as far back as their own premises.

Each of these boundaries may be called an edge, and depending on the requirements of the use case, a different set of access technologies or interfaces are used to fulfil the same end goal of bringing computing and storage closer to the user.

Edge computing may also be seen as a set of technologies that aims to bring distributed cloud computing closer to the user. The Linux Foundation (LF) edge consortium has been a key player in advancing the architecture of edge computing [LFEdge2020][63].

Figure 18 illustrates the edge 'computing continuum', which traverses from constrained and highly distributed user edge to centralised data centres while crossing multiple edge boundaries from user edge to service provider edge. Each of the boundaries includes design trade-offs that system designers need to bring compute resources closer to the physical world.

The LF edge taxonomy can be viewed as an illustration of the future Internet data transportation infrastructure. The ecosystem is expected to evolve towards ubiquitous compute fabric, allowing applications running over the distributed infrastructure to optimise the compute and storage location based on latency, bandwidth, autonomy, security and privacy, cost and energy efficiency.

[63] For edge computing glossary of terms, see [LFEdge2020].

Figure 18: Sharpening the Edge: overview of the LF Edge taxonomy and framework. Source: CC [LFEdge2020].

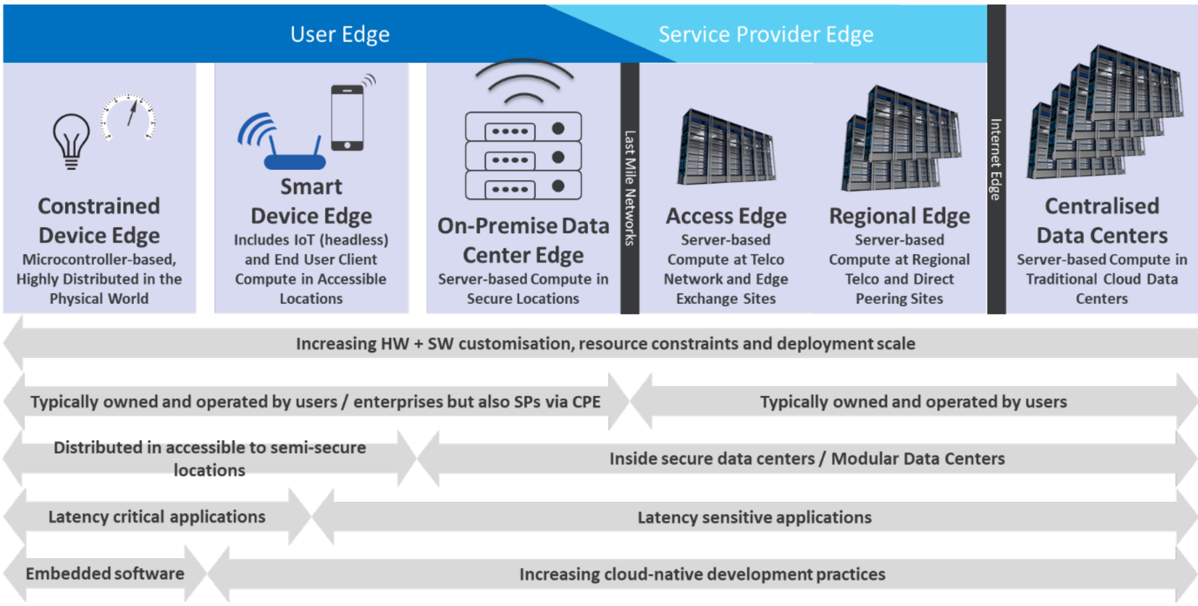

Closer edge to the end user

The edge may be even be pushed further towards the end user because of its ability to lower latency, lessen the network load/cost, higher energy efficiency, and better privacy.

For example, it is reasonable to imagine that mobile devices such as smartphones could act as the edge. In this scenario, a smart watch or a health sensor would upload its data onto the smartphone over a personal or body area network, where the data can be stored or processed before pushing it further towards the cloud core.

Similarly, the customer premise equipment (CPE) could be evolved within the context of wider notion of convergence to provide compute and storage capability as well as connectivity. In the latter case, multiple connectivity on the wide-area-network (WAN) side[64] and similarly on the local-area-network (LAN)[65] side. On the WAN-side, fixed wireless access (FWA) or cellular connectivity may be added to current CPE thus allowing a particular premise to gain access to both types of WAN technologies. Similarly, on the LAN-side, Wi-Fi and Ethernet could be added to future femtocell gateways, see Figure 19.

Figure 19: Potential future CPE. Source: Ofcom.

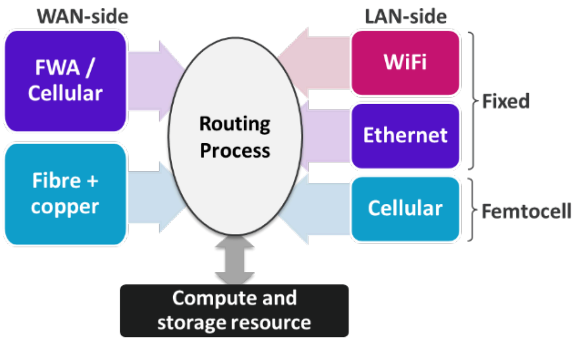

In this way and from the cloud-edge perspective, the smarter CPEs allow in-premise compute and storage thus covering consumer requirements such as privacy. From connectivity perspective, the CPE architecture increases WAN bandwidth with reliability. It also allows for guaranteeing QoS via the femtocell while maintaining interoperability with a vast range of devices via Wi-Fi. Finally, this architecture will allow exploitation of

[64] The broadband provider facing side.

[65] The consumer facing side.

multi-path TCP as part of QUIC to maximise the end-to-end throughput which leads to greater consumer QoE.

Shift towards open modularisation

Historically, routing and switching devices have been limited to tightly coupled hardware and software components. More recently, a 'disaggregated' approach has been developing which decouples the software component from the hardware. With this, white-box network switching hardware have emerged along with Network Operating Systems (NOSs) that support multiple hardware vendors. The hardware-software disaggregation in white box switches has pushed open-source NOS to be developed.

Software for Open Networking in the Cloud, alternatively referred to as SONiC, is a Linux-based open source NOS [SONiC][66]. It offers a full suite of network functionality, like BGP and Remote Direct Memory Access, and runs on switches from multiple vendors. It consists of several modules that exist either in containers (see next section) or in the Linux-host system itself. Open source NOSs and white-box switches are targeted to large data centres, and they may also be used in campus networks.

The benefits of disaggregated network devices are the potential for lowering cost and faster refreshing of software, be it at the NOS or container level.

Evolution in cloud software environment

The inception of virtualisation started the mass migration to the cloud and made clouds feasible. Hypervisor virtualisation technology allowed the running of multiple virtualised servers, or virtual machines (VMs), concurrently on one physical box [KingEtAl2003]. Next, container technology emerged which bundles of one or more applications and the software dependencies needed for that application to run into one highly portable package. When compared with VMs, a container footprint is much smaller, and needs a minimal amount of Operating System (OS) memory[67] [CloudHedge2020]. More recent additions to the family of virtualisation technology include:

- MicroVM which are hardware-isolated lightweight VMs with their own mini-kernel [Jain2020]. Besides being efficient and high-performance, microVMs can be used to protect computing devices against the execution of malicious code and similarly protect applications and data running on untrusted machines.
- Unikernels which are tiny VMs such that extremely streamlined/minimalistic OS is linked directly with the target application [GoethalsEtAl2018]. Unikernels are very quick to start because of the drastic reduction in the OS and this makes them viable for on-demand services.

WebAssembly: A potential future addition to the aforementioned solutions is WebAssembly for enabling low latency functions and the usage of limited resources. It was discussed in chapter 2 in the context of client-side and web browsers. More recently, there has been a realisation of WebAssembly's potential to transform software development beyond the browser by overcoming the security and

[66] Note that this is a different SONiC than the Sonic mentioned in chapter 1, and distinct from Ofcom's SONIC Labs facility [OfcomSONIC2021].

[67] A container could be tens of megabytes in size, whereas a VM may be several gigabytes in size. Hence, a single server can host far more containers than VMs.

stability challenges of relying on untrusted software from third parties.

The Bytecode Alliance has been formed by industry players to further advance WebAssembly beyond the browser through innovations in compilers, runtimes, and tooling, fine-grained sandboxing, capabilities-based security, modularity, and standards with security, efficiency, modularity, and portability in mind [BytecodeAlliance].

WebAssembly is a nascent technology that provides a strong memory isolation (through sandboxing) at near-native performance with a small memory footprint. There has been a significant effort in the last few years towards adopting WebAssembly for native execution because of its portability and potentials support for different high-level languages such as C/C++, Rust and Go. Figure 20 gives an overview of server-side WebAssembly [SalimEtAl2019].

Figure 20: Overview of server-side WebAssembly. Source: Ofcom.

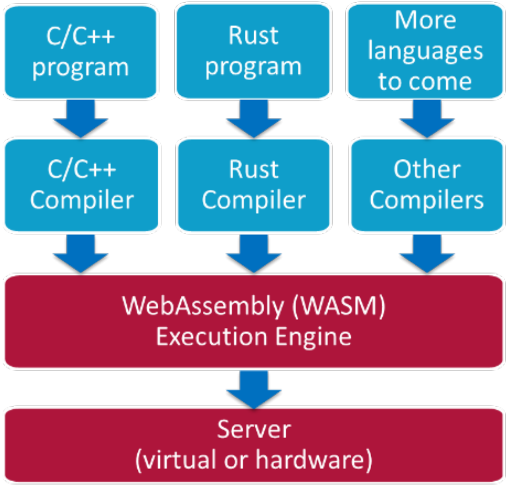

4: The Internet and energy efficiency

This chapter reports on some of the opportunities to reduce energy consumption in tomorrow's Internet, specifically in cloud and edge infrastructure which spans across elements of data centre, mobile network, and fixed network operators. While it is important to note that efforts are being driven globally to reduce the carbon emissions across the entirety of the information and communications technology (ICT) sector, including in areas such as consumer devices and network equipment, this chapter considers how future cloud-edge infrastructure may be built and operated with environmental sustainability in mind.

Context

Currently, the ICT sector represents between 1.5% to 4% of global greenhouse gas emissions, and 3% of global energy [Allard2020]. A recent study has found that efficiency improvements and the use of renewable energy sources have kept ICT's carbon footprint relatively flat in recent years [Malmodin&Lundén2018]. That said, meeting the 1.5°C Paris agreement target will require a reduction in carbon emissions though efficiency improvements, longer device lifetimes and higher use of renewable energy.

In a recent joint paper, the ITU, GeSI, GSMA and the Science Based Targets Initiative described a set of greenhouse gas (GHGs) reduction trajectories for ICT operators which are consistent with the effort to limit global warming to 1.5°C [GSMA&ITUEtal2020, ITU2020]. Figure 21 summarises the year-on-year GHG emission reduction trajectories for data centre, mobile network, and fixed network operators.

Figure 21: Greenhouse gas emission Trajectories for ICT operators with percent reductions from 2020 to 2030. Source: [GSMA&ITUEtal2020].

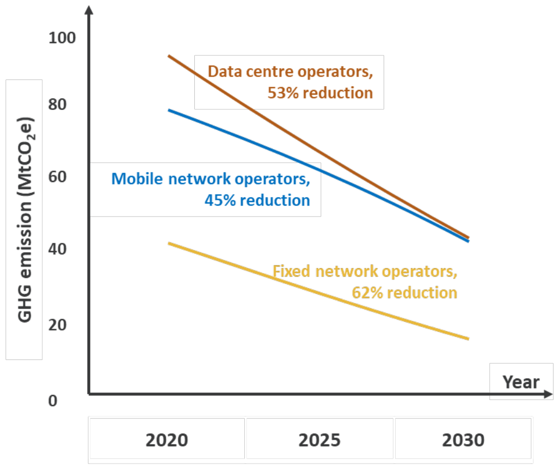

The opportunities described in this chapter could enable ICT operators to reduce energy consumption in their overall Internet cloud and edge infrastructure, which could then help to:

1) Reduce the dependency on non-renewable energy, and/or
2) Increase the efficiency of use of renewable energy.

Together, they help increase the likelihood of reaching the target of 'net zero' by 2050 [GSMA&ITUEtal2020].

Reducing energy consumption in cloud-edge infrastructure lessens the reliance on nonrenewables in the interim period prior to 2050. In the long term, it saves renewable energy thus making it available for use in other areas and reducing operational costs.

Opportunities to reduce energy consumption could come from the three areas listed below, which are broadly addressed in this chapter:

1) Technology that is used to build cloud-edge infrastructure including processor and storage technologies.
2) Where cloud-edge infrastructure is built[68], how the server farms are constructed and how the infrastructure is inter-networked.
3) How cloud-edge infrastructure is operated to optimise for energy consumption and QoS.

Figure 22: Architecture of cloud-edge infrastructure. Source: CC [AlharbiEtAl2020].

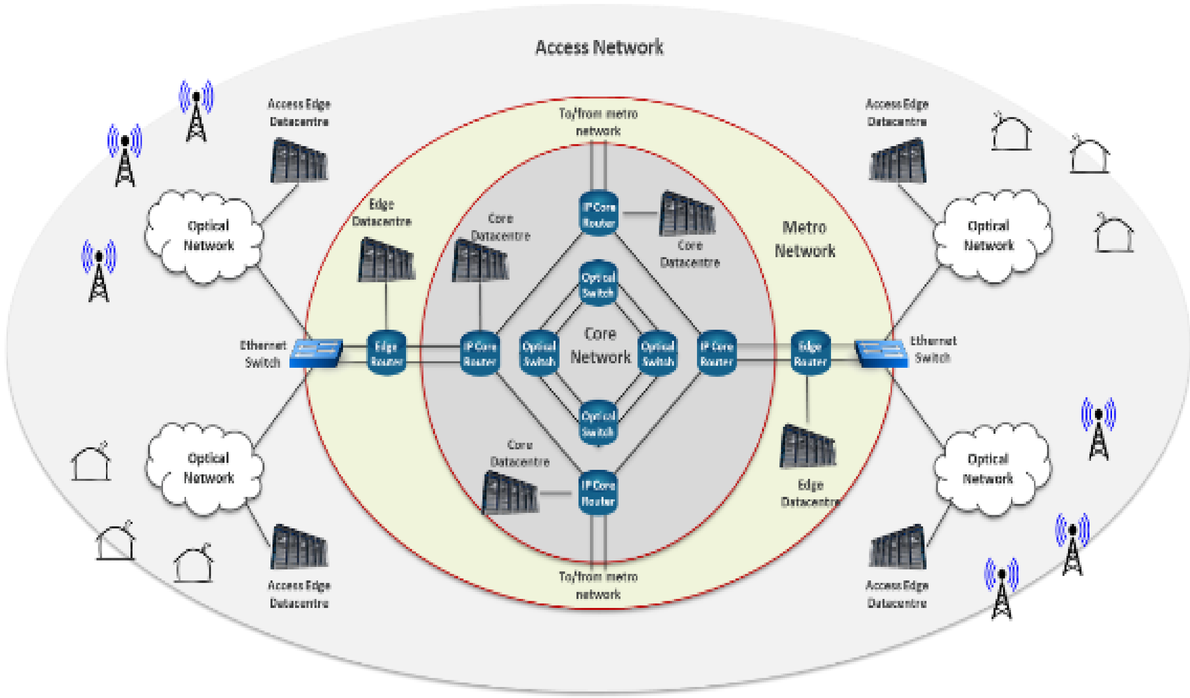

Architecture of future cloud-edge infrastructure

Figure 22 shows a modern cloud infrastructure which serves both mobile and fixed access networks. This architecture is regarded as an effective approach to deliver future services because of the interplay between edge and core cloud infrastructures[69].

The infrastructure includes edge data centres placed in the access and metro network; and data centres in the core network [AlharbiEtAl2020].

---

[68] Given that electrical energy transmission losses outweigh the energy lost in transmitting data through communication networks [DongEtAl2011].

[69] It is noted that the edge is depicted to be in the access network. In contrast, the work in chapter 3 indicated that the edge is related to the point of service, which can be the mobile device and the CPE in the case of mobile and fixed network, respectively. This point can be in the core cloud through to the end device. Thus, the edge concept in Figure 22 may be taken as a 'hard' form of edge while the dynamism implied in the work in chapter 3 can be understood as a 'soft' form of edge.

Conventional data centres operated by a cloud service provider are part of a large network infrastructure spanning across the world. The data centres are usually interconnected through a dedicated backbone network or through existing core networks over which services are delivered to end users with data being migrated between the data centres for optimality. Access networks are also connected to data centres in the edge, metro and core.

The data centres comprise hundreds of thousands of conventional servers and are supported by network management; and energy, cooling and security systems to name a few.

The general-purpose nature of this infrastructure allows it to be used for various applications. The following is a non-exhaustive list of applications:

- The provisioning of software defined networks, compute and storage to meet enterprise requirements;
- The provisioning of network function virtualisation for future mobile networks;
- The provisioning a platform for content distribution; and
- The provisioning of a platform for IoT and smart cities.

Some of these applications are discussed later in the chapter.

From a sustainability perspective, optimal use of this distributed cloud-edge infrastructure is key to further reducing ICT sector contribution to global emissions. Cooperation between the edge and the core cloud is the vehicle we focus on in this chapter for more efficient and greener computing platform [HuEtAl2017, DengEtAl2015].

Energy savings due to better technology

This section presents a number of advances in the way data centres are built which could help to reduce energy consumption.

Battery capacity and processor technologies are improving

One key role of data centre operators is to strive for higher levels of availability of service (up time). Currently, they employ fossil-based diesel generators for backup and thus mitigate the risk of interruption of normal energy supply, which is generally based on a mixture of renewable and non-renewable sources [Patrizio2021b].

Battery technology helps in two ways. First, it reduces the reliance on non-renewable energy. Second, it reduces the need to transport energy between the generation and consumption points.

Lithium-ion batteries are being used in data centres because of their better energy density compared to lead-acid batteries [Poe, Brown2019][70]. Current state-of-the-art lithium-ion batteries have an energy density of 240-270 Wh/kg [Faraday2020][71].

A lithium-metal battery is considered to be superior in battery chemistry because of its high capacity and energy density [Burrows2021]. Recently, a stable proof-of-concept lithium-metal solid state battery has been designed that has greater longevity and which, due to its high current density, could be fully charged within 10 to 20 minutes

[70] Lithium-ion battery technology has been used in electric vehicles, electric aircraft, unmanned air (drones), maritime, and land vehicles.

[71] Tesla Model 3 uses an estimated 250 Wh/kg Lithium-ion battery made by Panasonic [Hawkins2019].

[Ye&Li2021]. The combination of high density reported up to 651 Wh/kg and fast re-charge time means such technology could help in harvesting intermittent renewable energy sources and serving ICT equipment over long periods where such sources of energy were absent.

Processor technology also plays a role from an energy efficiency perspective. As discussed in Chapter 3, recent advances in 2nm process technology suggest that up to 75% less energy can be achieved compared to the 5nm chip technology [Gartenberg2021, Patrizio2021a]. Similarly, the co-packaged optics connect technology discussed in Chapter 3 is also reported to use less energy than previous generations [Gartenberg2021, Patrizio2021a].

Energy savings due to where and how data centres are built

This section presents a number of advances in the location of data centres and how they are built to help to reduce energy consumption.

Proximity of data centre to energy sources lowers losses due to energy transmission

The availability of renewable energy resources is another important factor to be considered when placing data centres. A study examined whether it is preferable to locate data centres near renewable energy sources (which are typically far from cities), or to locate data centres near to population centres and transmit renewable energy to these data centres [DongEtAl2011]. The study found that locating the data centres near renewable energy sources was optimal from the total energy consumption perspective, as it maximises the utilisation of renewable energy and reduces the electrical energy transmission losses, which far outweigh the energy lost in communication networks in transmitting data to remote data centres. For the network used in the study, a CO_2 reduction of up to 73% was achieved in comparison with the scenario when non-renewable energy is being used and data centre locations are non-optimal (i.e. far away from the renewable energy source).

The above cited study indicated the importance of the proximity of data centre to energy sources. While locating data centres near renewable energy sources can negatively impact delay-sensitive real-time applications, it is effective in reducing energy consumption for large data compute (simulation and modelling of complex systems), storage of long-tail content, batch processing of large amount of data, and long-term data storage (achieves). The delay-sensitive real-time applications are dealt with in the next section.

Proximity of data centre to population areas improves QoS

In another study, a hybrid approach was proposed to overcome the above described negative impact: access edge data centres placed near the end users and large data centres placed away from users [MohamedEtAl2020]. The edge cloud achieves optimum results from energy consumption and QoS. However, the limited compute and storage of mini-data centres maintains the need for large data centres which can be placed away from the population areas (not necessarily near energy sources), and hence offers some degree of energy saving through economies of scale in processing.

In this hybrid scenario, energy savings could be achieved with dynamic management of service provisioning and optimisation where compute and storage are dynamically provided based on energy consumption, QoS requirements and variable traffic volume. For the scenario studied, it was found up to 75% energy savings could be achieved in comparison with the case when no edge computing is available [MohamedEtAl2020].

It is noted that this approach can also lead to smaller/mini edge data centres in populated areas which increase the effectiveness of renewable energy usage as an alternative to electrical energy drawn from the national grid by the mini edge data centres.

Submerged data centres live longer and need less energy

Recent research addressed the concept of underwater data centres [Judge2021, Roach2020]. In one study, several hundreds of servers and over 27 petabytes of storage were tightly packed into a steel cylinder, filled with dry nitrogen[72]. This data centre was placed on the seafloor of the North Sea at a depth of more than a hundred feet.

The underwater data centre was found to be environmentally and economically practical. The submerged data centre was eight times more reliable than land data centres, because it is cocooned from corrosive oxygen and moisture. The submerged data centre had a power usage effectiveness (PUE)[73] ratio of 1.07, while the average PUE ratio for a land-based data centre in 2020 has been reported to be 1.58 [Lawrence2020]. This represents approximately 32% reduction in energy consumption and is attributed to natural seawater acting as the coolant instead of active energy-consuming methods [Christian2020] [74].

Disaggregated server farms need less energy

A conventional server (CS) is the basic computational unit in a conventional data centre. It comprises of dedicated compute, storage, and network resources. Up to about 48 CSs are typically arranged into a single rack cabinet in a conventional data centre as illustrated in Figure 23 [AjibolaEtAl2021].

Figure 23: A conventional server rack in a data centre[75]. Source: [AjibolaEtAl2021].

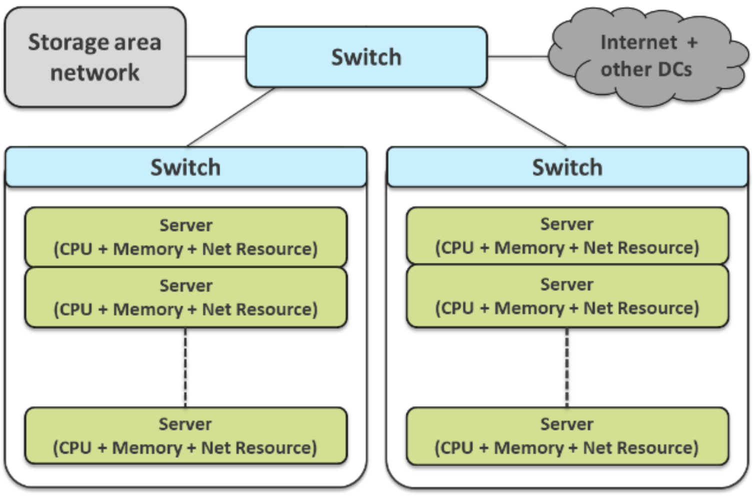

[72] The storage was sufficient to store nearly five million movies, and its compute capability was as powerful as several thousand high-end consumer PCs.

[73] Power usage effectiveness is a ratio that describes how much energy is used by the computing equipment in contrast to other overhead that supports the equipment (e.g. cooling, etc). The closer the PUE number to unity the higher the level of efficiency.

[74] Cooling is reported to consume approximately 20% of energy used by land data centres [Christian2020].

[75]

Disaggregated server architecture is a promising approach to achieve reductions in energy consumption. Here, the CS components are separated into individual physical or logical compute, storage, and network pools. The resources of each individual pool are orchestrated on-demand to form logical servers that support workload execution.

After successful execution, the orchestrated resources are returned to their respective pools and made available for the formation of new logical servers.

Physically, practical disaggregation of CS components can be achieved at rack-scale (RS) or at pod-scale (PS) as illustrated in Figure 24.

Figure 24: Rack-scale and pod-scale disaggregated data centre[76]. Source: [AjibolaEtAl2021].

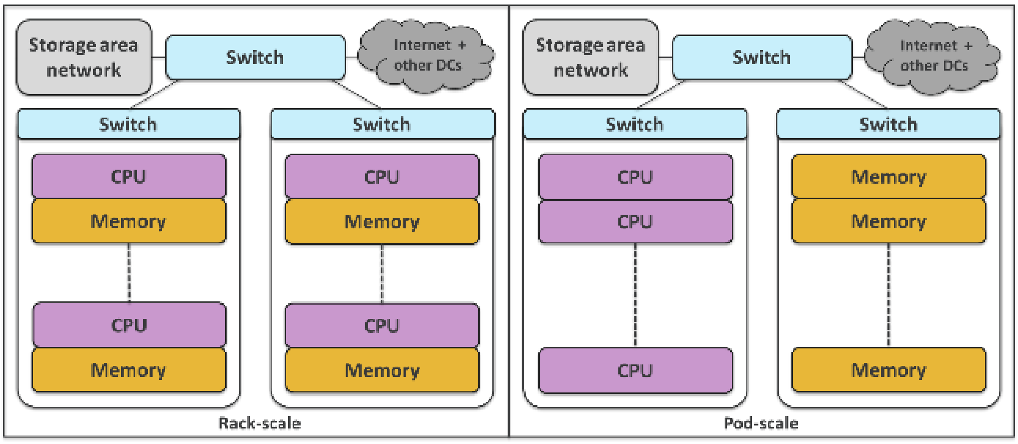

In the RS case, each server-like node contains different resource types which are placed in the same rack: each rack contains a mixture of compute, or alternatively known as Central Processing Unit (CPU), and memory units. These resources are selected to form a logical configuration server. In contrast and in the PS case, the same resource type is placed in the same rack: each rack contains either a set of computes (CPU) or memory units. These resources selected form a logical configuration server. This data centre design introduces significant efficiencies, but also introduces network challenges which were either non-existent or trivial in conventional data centres.

Three types of processing were studied: processing intensive, memory intensive, and input-output intensive [AliEtAl2017]. For these scenarios, the study found that disaggregated data centre achieves up to 42% average savings in total energy consumption relative to conventional data centres and similar energy savings can be achieved in the case of RS.

Mesh and star network architectures are optimal for core

Topology planning is fundamental to building network infrastructures not only from traffic profile, capacity, and latency perspectives, but also for energy consumption. Essentially, there are four network topologies that form the basis of many modern core networks namely: star, ring, line and mesh topologies.

The energy consumption of core networks has been the subject of recent research [MusaEtAl2018]. Figure 25 shows that mesh

[76] © 2021 IEEE. Reprinted with permission from [AjibolaEtAl2021].

topologies having lower energy consumption than other topologies, star topology has a close but higher consumption.

Figure 25: Theoretical energy consumption of different network topologies[77]. CC Source: [MusaEtAl2018].

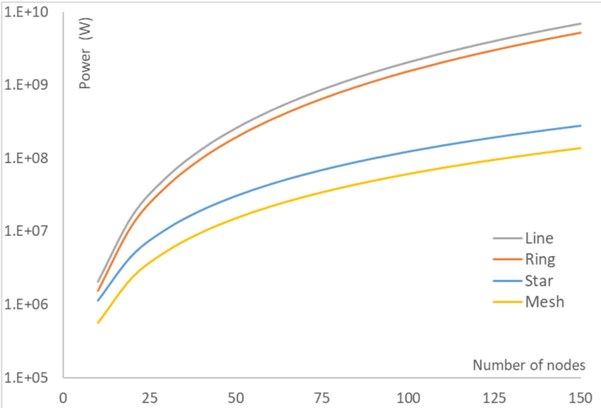

Line topology consumes more energy than the ring topology, which in turn has a higher consumption than the star. The plot is for a typical European sized country whose geographic distances are in the order of 100km and with up to 150 nodes.

Energy efficiency due to network operation

Energy consumption is impacted by a number of operational factors which include the location where data is processed, stored and served; and the distance over which a volume of data needs to be transported.

Virtual machine placement in cloud-edge networks with maximal efficiency

An effective optimisation scheme for resource allocation is crucial for the cloud infrastructure shown in Figure 22 because services can be placed in a highly energy inefficient server or even further from the source node or consumer, resulting in higher communication latency. It is expected that greater energy efficiency can be achieved through optimal distribution of virtual machines (VMs) between the access and metro edge nodes and the core nodes [HuEtAl2017].

It has been shown that optimum placement of VMs that takes into consideration energy consumption over this computing continuum could achieve substantial energy savings, between 56% and 64% under high user data rates, compared to the non-optimised VMs distribution [AlharbiEtAl2020].

Software defined networks in data centres

The importance of Software Defined Networking and Network Function Virtualisation (NFV) have been long established as pillars of future networks that promise to support applications such as corporate services, enhanced mobile broadband, ultra-low latency and massive sensing type applications while providing resiliency in the network [ETSI2012, Dutta2020]. They are built on a virtualisation-based infrastructure across data centres.

Cloud users submit their requirements for a software defined network and compute resource to a cloud service provider. Each request consists of the specification of multiple virtual nodes (compute, storage, and networking resources), all of which are to be interconnected through a set of virtual links. The cloud service provider maps the requested virtualised resources onto their physical network infrastructure. When successfully mapped, a 'virtual slice' consisting of the requested virtual nodes will have been established across the geographically distributed cloud-edge infrastructure and the

[77] This is for a very large high-capacity network. Note that the ring curve is for even number of nodes, which is comparable to ring networks with odd number of nodes.

requester is granted ownership of the virtual slice[78].

The mapping may be done in two different ways: the 'cost of available bandwidth' approach or the 'energy aware approach', where the former seeks to minimise the number of activated nodes and links without accounting for the energy consumption of the data centres.

Optimal mapping and provisioning have been shown to achieve energy savings of up to a maximum of 60% (20% on average) in an environment where cloud user requests are being dynamically allocated and de-allocated [DongEtAl2011].

Network function virtualisation and future mobile networks

A typical legacy cellular network uses a dedicated set of hardware for every network function. In this way, energy inefficiencies occur because of the potential for hardware to have low utilisation in some sites. Figure 26 shows the role played by NFV in a cellular network, which paves the way for optimising the network resources from an energy efficiency perspective [Al-QuzweeniEtAl2019].

Here, an optical network backbone is used to support the mobile 5G network. Thus, the mobile core network functions, together with baseband functions, are virtualised and abstracted in the form of VMs. Figure 27 shows the variation of energy consumption in a 24-hour daily cycle [Al-QuzweeniEtAl2019].

Figure 26: Cellular networks with network function virtualisation. Source: CC [Al-QuzweeniEtAl2019].

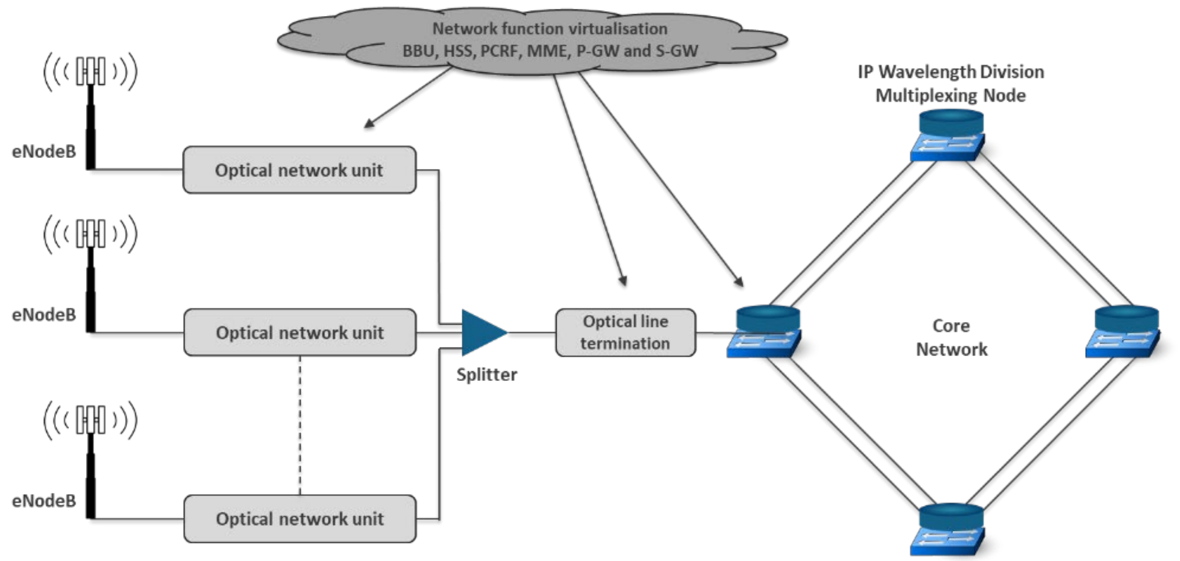

The study found that in general non-virtualised infrastructure consumes larger amount of energy because it is incapable of moving the processing load into few machines during low-load periods. The results indicated that energy savings of up to 38% (34% on average) are possible with optimised NFV when energy efficiency algorithms are used to optimise VMs location in view of improving server utilisation with respect to network loading, with the goal

[78] A slice is a logical interlinking of virtual compute and networking resources that have been established to meet a particular demand.

of minimisation of overall energy consumption and maximising QoE.

Figure 27: Comparison of energy consumption the cellular network with (VIR) and without virtualisation for different time of the day and under different inter-virtual machine traffic. Source: CC [Al-QuzweeniEtAl2019].

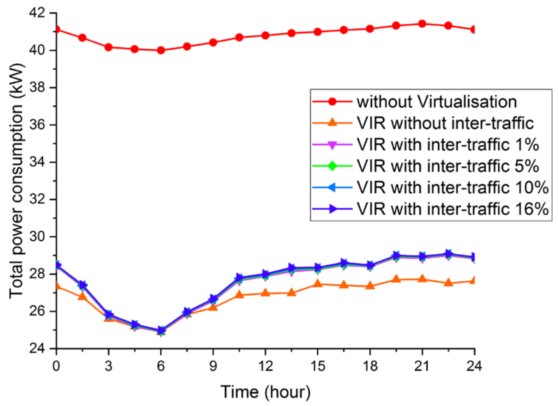

Content distribution and caching

Content Distribution Networks (CDNs) are geographically distributed set of servers which work together to provide fast delivery of Internet content by storing frequently used content at locations close to the consumers of that content [79].

Traditionally, CDNs have been placed in the Internet core. Recent trends have pushed CDNs into access networks to reduce the CDN-to-consumer distance and improve QoE. With video distribution contributing greatly to the Internet traffic, CDNs can consume a large amount of energy.

Thus, within the context of Figure 28, CDNs can be placed in various points within the geographically distributed cloud-edge infrastructure [OsmanEtAl2014]. Efficiency can be achieved with real-time dynamic optimisation of the location of CDN provisioning with respect to load and energy consumption.

Figure 28: System model for CDN/cache network over WDM, metro and access networks[80]. Source: [OsmanEtAl2014].

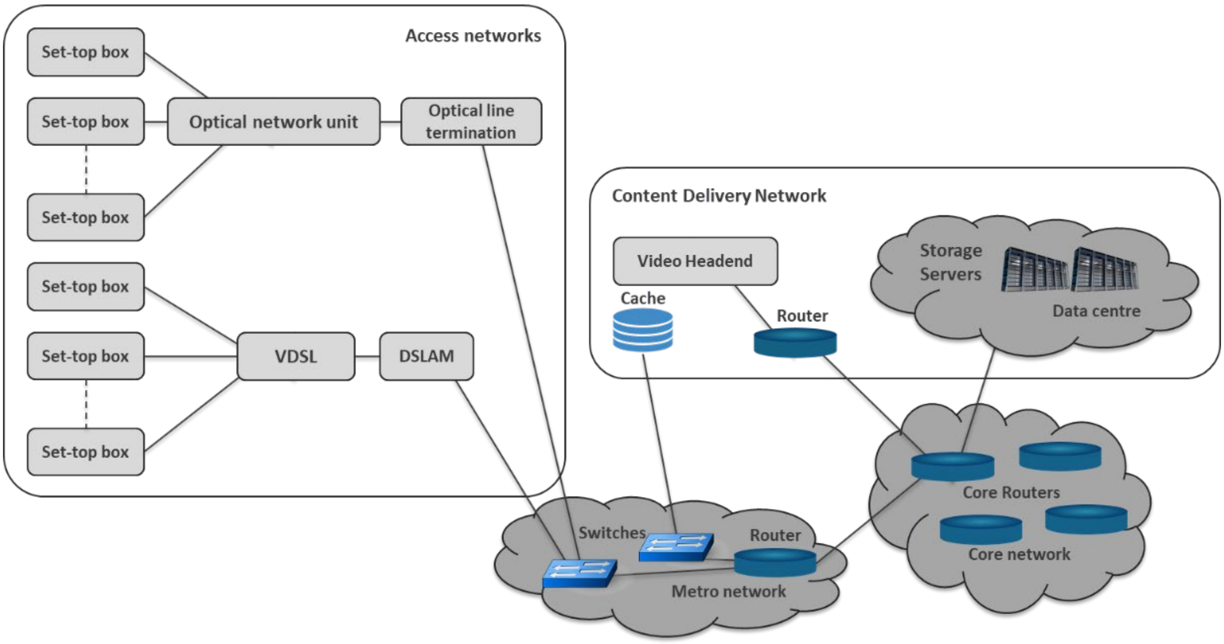

Some studies have shown that storing popular content towards the edge achieves an instantaneous power reduction of up to 36% including caching for the network of Figure 28,

[79]They play an important role in the video distribution from vendors such as Netflix, Amazon and Facebook.

[80] © 2021 IEEE. Reprinted with permission from [OsmanEtAl2014].

whose total energy consumption of can be reduced by up to 86% with variable size caches for the considered scenario in the study as compared to no caching [OsmanEtAl2014].

In a similar study to the above, reduction in energy consumption can be achieved from the replication of content based on popularity in multiple clouds under energy optimisation and control [LaweyEtAl2014a]. The cited study indicated up to 43% total energy saving could be obtained in the studied scenario when compared with the equivalent case where content stored in one CDN infrastructure and with no energy management.

In addition, peer-to-peer (P2P) and hybrid P2P-CDN architectures has been suggested to lead to energy consumption savings of 61% and 32% respectively, compared to single CDN with no energy management [LaweyEtAl2014b].

Efficient virtualised Internet of things

As the number of IoT devices increases, the data volumes generated by these end devices will also increase and so will their energy consumption. Thus, even though in some cases such networks may offset energy savings by other sectors, there is still a need to address the energy consumption aspects of the resultant large IoT networks.

An energy efficient IoT virtualisation framework that supports P2P networking and processes at the edge of the network is illustrated in Figure 29, where the processing tasks can be served in VMs that can be hosted by processors in IoT objects and/or in processors in relay devices [Al-AzezEtAl2019].

This leads to three scenarios in terms of where data is processed: in relay devices, in IoT objects, and a hybrid scenario where data processing is done in both relays and IoT objects. Up to 62% energy saving is possible when optimising across the three scenarios [Al-AzezEtAl2019].

Figure 29: Energy efficient IoT architecture with P2P communication and processing[81]. Source: [Al-AzezEtAl2019].

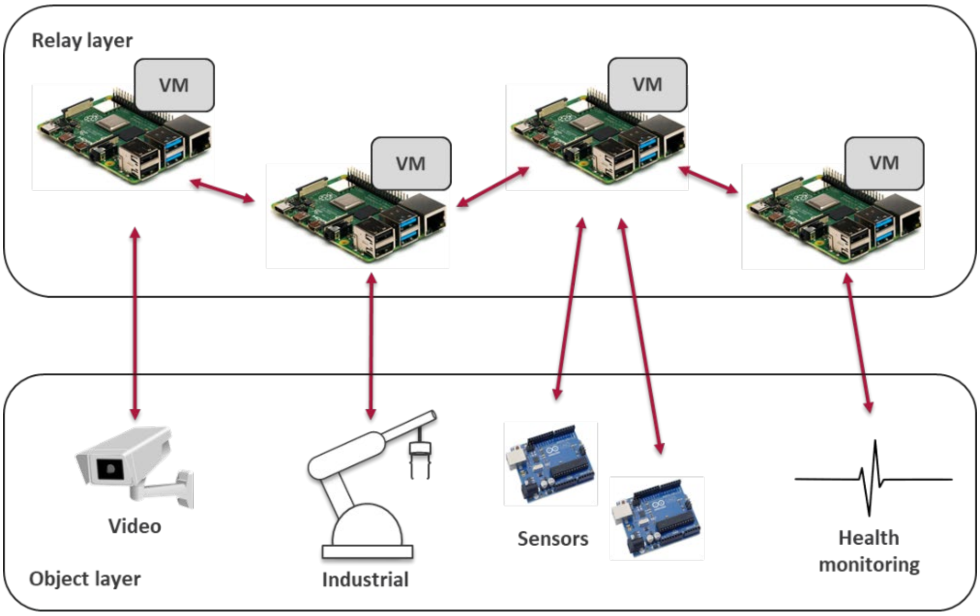

[81] © 2021 IEEE. Reprinted with permission from [Al-AzezEtAl2019].

IoT in smart cities: To manage large IoT networks which span across smart buildings, smart factories, smart homes, etc, a service oriented approach has been proposed [Abbas&Yoon2015, YousefpourEtAl2019]. As highlighted in these studies, the benefit of this approach has been said to make sensor information available on-demand thus establishing and tearing down services as required in rea time.

One study looked at energy efficiency gained from minimising network and processing energy consumption; minimising mean traffic latency on its own and minimising a weighted combination of both the total energy consumption and traffic latency [Al-ShammariEtAl2019]. Compared to the baseline reference where demand requirements are met with no consideration to energy consumption or latency, the study reported a total energy consumption saving of up to 42% can be achieved and an average traffic latency reduction of 47% obtained when 'local' edge processing is available. The savings can be as high as 86%, when compared to all processing taking place in the core cloud [YosufEtAl2020].

References

R.1 Foreword

[Ofcom2021] The Office of Communications, 'Technology Futures – spotlight on the technologies shaping communications for the future', Jan. 2021.

R.2 Summary

[Brahimi2005] L. Brahimi, 'Conversations with History', Interview by the Institute of International Studies, University of California, Apr. 2005.

[DeNardis2020] L. DeNardis, 'The Internet in everything: Freedom and security in a world with no off switch'. Yale University Press, Jan. 2020.

[Hernández-MuñozetaletalEtAl2011] J. M. Hernández-Muñozetal et al, 'Smart cities at the forefront of the future Internet', pp. 447-462, in J. Domingue et al, The Future Internet Assembly, Lecture Notes in Computer Science, Springer, Jan. 2011.

[IS2019] Internet Society, 'Consolidation in the Internet Economy', Global Internet Report, 2019.

[ITU-Stats] International Telecommunications Union, 'Number of Internet users at the end of 2019 globally', accessed online in Jun. 2021.

[Kamiya2020] G. Kamiya, 'Data centres and data transmission networks', International Energy Agency Tracking Report, Jun. 2020.

[McLuhan1962] M. McLuhan, 'The Gutenberg Galaxy: The Making of Typographic Man'. University of Toronto Press, Jan. 1962.

[O'Hara&Hall2021] K. O'Hara and W. Hall, 'Four Internets: Data, Geopolitics and the Governance of Cyberspace'. Oxford University Press, Aug. (forthcoming) 2021.

[Ofcom2021] The Office of Communications, 'Technology Futures – spotlight on the technologies shaping communications for the future', Jan. 2021.

[RenEtAl2019] J. Ren, D. J. Dubois, D. Choffnes, A M. Mandalari, R. Kolcun and H. Haddadi, 'Information exposure from consumer IoT devices: A multidimensional, network-informed measurement approach', Internet Measurement Conference, Oct. 2019.

[Stephenson1992] N. Stephenson, 'Snow Crash'. Random House Publishing Group, 1992.

R.3 Methodology

[Ofcom2020] The Office of Communications, 'Call for inputs on the technologies of tomorrow', Jul. 2020.

[ParzialeEtAl2006] L. Parziale, D. T. Britt, C. Davis, J. Forrester, W. Liu, C. Matthews and N. Rosselot, 'TCP/IP tutorial and technical overview'. IBM, Redbooks, Dec. 2006.

R.4 Everyday life with the future Internet

[5GACIA2019] 5G Alliance for Connected Industries and Automation, '5G for Connected Industries and Automation Second Edition', 5GACIA White Paper, Feb. 2019.

[AghaeiEtAl2012] S. Aghaei, M. A. Nematbakhsh and H. K. Farsani, 'Evolution of the World Wide Web: From Web 1.0 to Web 4.0', International Journal of Web & Semantic Technology, Jan. 2012.

[AjuntamentDeBarcelona] Ajuntament de Barcelona, 'Household waste collection', accessed online in May 2021.

[Almeida2017] F. Almeida, 'Concept and dimensions of Web 4.0', International Journal of Computers and Technology, Nov. 2017.

[Atwell2017] C. Atwell, 'Industry 5.0 is already on the horizon', Machine Design, Sep. 2017.

[Baker2015] A. Baker, 'Connected Spaces: The next step for the Internet of Things', The Guardian, Feb. 2015.

[Bluetooth] Bluetooth, 'About us', accessed online in Jul. 2021.

[Brügger&Schroeder2017] N. Brügger and R. Schroeder, 'The web as history: Using web archives to understand the past and the present'. UCL Press, Mar. 2017.

[Callan2021] T. Callan, 'What are machine identities and why should engineers care?', Electronic Products, Jun. 2021.

[Castro2021] C. Castro, 'Tasting Digital: How the Way you Sense the World Will Change in the Next Decade', 6G World, Jan. 2021.

[CBS2021] Team at CBS Miami, 'Scientists working on robot that can detect all kinds of emotions in hopes of helping patients with mental health disorders', Jun. 2021.

[Choudhury2014] N. Choudhury, 'World Wide Web and its journey from Web 1.0 to Web 4.0', International Journal of Computer Science and Information Technologies, 2014.

[CookEtAl2020] A. V. Cook et al, 'The spatial web and Web 3.0: What business leaders should know about the next era of computing', Deloitte, Jul. 2020.

[CornellU2018] Cornell University, 'Emotional robot lets you feel how it's 'feeling'', ScienceDaily, Jul. 2018.

[DataTransferProject] The Data Transfer Project, 'About us', accessed online in Jun. 2021.

[Deloitte2015] Deloitte, 'Industry 4.0: Challenges and solutions for the digital transformation and use of exponential technologies', 2015.

[DemirEtAl2019] K. A. Demir, G. Döven and B. Sezen 'Industry 5.0 and human-robot co-working', Procedia Computer Science, Jun. 2019.

[DeNardis2020] L. DeNardis, 'The Internet in everything: Freedom and security in a world with no off switch', Yale University Press, Jan. 2020.

[ENRD2018] European Network for Rural Development, 'Smart Villages Revitalising Rural Services', EU Rural Review No 26, May 2018.

[Ericsson2019] Ericsson, 'Ericsson ConsumerLab: Ten hot consumer trends 2030 – The Internet of Senses', Ericsson News, Dec. 2019.

[Ericsson2020] Ericsson, 'The dematerialized office: A vision of the Internet of Senses in the 2030 future workplace', Oct. 2020.

[FeldmannEtAl2021] A. Feldmann et al, 'A year in lockdown: How the waves of COVID-19 impact Internet traffic', Communications of the Association for Computing Machinery, Jul. 2021.

[ForgeEtAl2018] S. Forge, R. Horvitz, S. Lindmark, E. Bohlin and C. Blackman, 'Fixed and mobile convergence in Europe', European Commission Study Report, Jun. 2018.

[FWB2018] Flat World Business, 'Web 1.0 vs Web 2.0 vs Web 3.0 vs Web 4.0 vs Web 5.0 – A bird's eye on the evolution and definition', Apr. 2018.

[Getting2007] B. Getting, 'Basic definitions: Web 1.0, Web. 2.0, Web 3.0', Practical Ecommerce, Apr. 2007.

[Gibson1984] W. Gibson, 'Neuromancer'. Ace Science Fiction Special, 1984.

[Goendoer&Sharhan2018] S. Goendoer and S. Sharhan, 'Sgoendoer/sonic', GitHub, May 2018.

[Google] Google, 'Nose Beta', accessed online in Jul. 2021.

[Gujral2019] R. Gujral, 'The future of AI: How emotion AI is making robots smarter', Forbes, Jul. 2019.

[Guo2021] J. Guo, 'Digital twins are shaping future virtual worlds', Service Oriented Computing and Applications, Mar. 2021.

[Hall2020] L. Hall, 'Living the future: The technological family and the connected home', Internet Matters Project Report, Aug. 2020.

[HemphillEtAl2019] A. Hemphill, S. Rabas, K. Arnold and O. K, 'Avatars: Are we ready for our digital twins?', BBC News, October 2019.

[Hernández-MuñozetaletalEtAl2011] J. M. Hernández-Muñozetal et al, 'Smart cities at the forefront of the future Internet', pp. 447-462, in J. Domingue et al, The Future Internet Assembly, Lecture Notes in Computer Science, Springer, Jan. 2011.

[IEC2019] The International Electrotechnical Commission, 'Semantic interoperability: Challenges in the digital transformation age', IEC White Paper, Oct. 2019.

[IFTF2017] Institute For The Future, 'Distributed reflections: The future of blockchain identity', 2017.

[ISO/IEC24760] ISO/IEC, 'IT security and privacy — A framework for identity management — Part 1: Terminology and concepts', ISO/IEC Standard, May 2019.

[Jardine2020] J. Jardine, 'Industry 5.0: Top 3 things you need to know', Master Control, Jul. 2020.

[JayantEtAl2010] C. Jayant, C. Acuario, W. Johnson, J. Hollier and R. Ladner, 'V-braille: Haptic braille perception using a touch-screen and vibration on mobile phones', Proceedings of the International ACM SIGACCESS Conference on Computers and Accessibility, Oct. 2010.

[Komninos2013] N. Komninos, 'Smart cities and the future Internet: Innovation ecosystems of embedded spatial intelligence', Proceedings of International Conference for Entrepreneurship, Innovation and Regional Development, Jun. 2013.

[Kwok2021] I. Kwok, 'Virtual reality project provides online space for architecture students', The Daily Californian, Mar. 2021.

[LaclavikEtAl2006] M. Laclavik, M. Babik, Z. Balogh, E. Gatial and L. Hluchy, 'Semantic knowledge model and architecture for agents in discrete environments', Proceedings of the European Conference on Artificial Intelligence, May 2006.

[LandTransportAuthority] Land Transport Authority, 'Intelligent transport systems', accessed online in May 2021.

[López2020] M. A. López, 'Self-sovereign identity: The future of identity: Self-sovereignity, digital wallets, and blockchain', LACChain Global Alliance digital identity working group, Sep. 2020.

[Marr2018] B. Marr, 'What is Industry 4.0? Here's a super easy explanation for anyone', Forbes, Sep. 2018.

[McClelland2017] C. McClelland, 'Do you know how haptics work?', Medium, Nov. 2017.

[Moore2018] S. Moore, '13 Surprising uses for emotional AI technology', Gartner, Sep. 2018.

[NathEtAl2015] K. Nath, R. Iswary and P. Borah, 'What comes after Web 3.0? Web 4.0 and the future', International Conference on Computing and Communication Systems, Apr. 2015.

[Nicolás2021] E. S. Nicolás, 'First glimpse of EU's new 'digital identity wallet'', EU Observer, Jun. 2021.

[O'Hara&Hall2021] K. O'Hara and W. Hall, 'Four Internets: Data, Geopolitics and the Governance of Cyberspace'. Oxford University Press, Aug. (forthcoming) 2021.

[Ofcom2021] The Office of Communications, 'Technology Futures – spotlight on the technologies shaping communications for the future', Jan. 2021.

[OMASpecWorks] Open Mobile Alliance SpecWorks, 'IPSO smart objects working group', accessed online in Jul. 2021.

[oneM2M] oneM2M, 'Who we are', accessed online in Jul. 2021.

[OpenConnectivityFoundation] Open Connectivity Foundation, 'Specifications', accessed online in Jul. 2021.

[Østergaard2017] E. Østergaard, 'Collaborative robots ushering in Industry 5.0', Universal Robots, Oct. 2017.

[OUK2019] OnUpKeep, 'What is the difference between Industry 3.0 and Industry 4.0?', May 2019.

[Özdemir&Hekim2018] V. Özdemir and N. Hekim, 'Birth of Industry 5.0: Making sense of big data with artificial intelligence, "The Internet of Things" and next-generation technology policy', OMICS: A Journal of Integrative Biology, Jan. 2018.

[PandoraFMS2018] Pandora FMS, 'What is Web 4.0?', Mar. 2018.

[PapadopoulosEtAl2021] G. Papadopoulos, P. Thubert, F. Theoleyre and C. J. Bernardos, 'RAW use cases' (work-in-progress), IETF Standard, Aug. 2021.

[PetitEtAl2019] O. Petit, C. Velasco and C. Spence, 'Digital sensory marketing: Integrating new technologies into multisensory online experience', Journal of Interactive Marketing, Feb. 2019.

[PwC2019] PwC, 'Blockchain and digital identity: The path to self-sovereign identity', 2019.

[Rene2019] G. Rene, 'An introduction to the spatial web', Medium, Oct. 2019.

[Riley2019] C. Riley, 'A framework for forward-looking tech competition policy', Mozilla, Sep. 2019.

[Saracco2019] R. Saracco, 'Digital twins: Bridging physical space and cyberspace', IEEE Computer Society, Dec. 2019.

[SchaffersEtAl2011] H. Schaffers, N. Komninos, M. Pallot, B. Trousse, M. Nilsson and A. Oliveira, 'Smart cities and the future Internet: Towards cooperation frameworks for open innovation', in J. Domingue et al, The Future Internet Assembly, Lecture Notes in Computer Science, Springer, Jan. 2011.

[Schwartz2019] O. Schwartz, 'Don't look now: Why you should be worried about machines reading your emotions', The Guardian, Mar. 2019.

[SofiaEtAl2021] R. Sofia, M. Kovatsch and P. Mendes, 'Requirements for reliable wireless industrial services' (work-in-progress), IETF Information, Mar. 2021.

[SONIC] SONIC, 'About', accessed online in Mar. 2021.

[Statt2021] N. Statt, 'Epic's new MetaHuman tool lets you craft realistic faces inside a browser', The Verge, Feb. 2021.

[Steiner1993] P. Steiner, 'On the Internet, nobody knows you're a dog', The New Yorker Cartoon, Jul. 1993.

[Tao&Qi2019] F. Tao and Q. Qi, 'Make more digital twins', Nature, Sep. 2019.

[TaoEtAl2019] F. Tao, H. Zhang, A. Liu and A. Y. C. Nee, 'Digital twin in industry: State-of-the-art', IEEE Transactions on Industrial Informatics, Apr. 2019.

[TheRoyalSociety2020] The Royal Society, 'RAMP update', Apr. 2020.

[Thilmany2007] J. Thilmany, 'The emotional robot. Cognitive computing and the quest for artificial intelligence', EMBO Reports, Nov. 2007.

[Weber&Rech2009] S. Weber and J. Rech, 'An overview and differentiation of the evolutionary steps of the Web X.Y movement: The Web before and beyond 2.0', pp. 12-39, in S. Murugesan, Handbook of research on Web 2.0, 3.0, and X.0: Technologies, business, and social applications. IGI Global, Oct. 2009.

[Webster2020] A. Webster, 'Travis Scott's first Fortnite concert was surreal and spectacular', The Verge, Apr. 2020.

[Widell&Keränen2020] N. Widell and A. Keränen, 'Data interoperability across IoT ecosystems with One Data Model (OneDM)', Ericsson, Sep. 2020.

[Windley2018] P. Windley, 'How blockchain makes self-sovereign identities possible', Computer World, Jan. 2018.

[Wong2021] Q. Wong, 'Twitter CEO Jack Dorsey explains his vision for social media's future', CNET, Feb. 2021.

[Wooldridge&Jennings1995] M. Wooldridge and N. R. Jennings, 'Intelligent agents: Theory and practice', The Knowledge Engineering Review, Jan. 1995.

[WB2021] The World Bank, 'Smart villages in Azerbaijan: A framework for analysis and roadmap', Mar. 2021.

[ZigbeeA] Zigbee Alliance, 'What is Zigbee?, accessed online in Jun. 2021.

[ZigbeeB] Zigbee Alliance, 'Dotdot', accessed online in Jun. 2021.

R.5 Evolution in the Internet and web technology

[Antipolis2020] S. Antipolis, 'ETSI launches new group on Non-IP Networking addressing 5G new services', ETSI, Apr. 2020.

[BeurdoucheEtAl2019] B. Beurdouche, K. Cohn-Gordon and R. Robert, 'Messaging layer security: Towards a new era of secure messaging...', BlackHat USA, 2019.

[BW2021] Business Wire, 'The connectivity standards alliance unveils matter, formerly known as project CHIP', May 2021.

[CardwellEtAl2021] N. Cardwell et al, 'BBR updates: Internal deployment, code, draft plans', IETF Presentation, Mar. 2021.

[Choffnes2016] D. Choffnes, 'A case for Personal Virtual Networks [PVN]', Proceedings of the ACM Workshop on Hot Topics in Networks, Nov. 2016.

[De-Ruiter&Schutijser2020] J. de Ruiter and C. Schutijser, 'New Internet infrastructures: An introduction to SCION', SIDN Labs, Sep. 2020.

[Duarte2008] R. P. Duarte, 'Transport protocols for large bandwidth-delay product networks - TCP extensions and alternative transport protocols', JETC Workshop on Electronics, Telecommunications and Computers Engineering, 2008.

[Durand2020] A. Durand, 'Resource Public Key Infrastructure (RPKI) technical analysis', ICANN, Sep. 2020.

[DW2020] Digitalisation World, 'Cross-sector consortium deliver data access platform', Jul. 2020.

[ETSI-NIN] ETSI, 'Non-IP Networking (NIN)', accessed online in Jun. 2021.

[ETSI2021] ETSI, 'Flexilink network model', Mar. 2021.

[FilsfilsEtAl2015] C. Filsfils, N. K. Nainar, C. Pignataro, J. C. Cardona and P. Francois, 'The segment routing architecture', IEEE Global Communications Conference, Dec. 2015.

[Galis&Lou2021] A. Galis and D.Lou, 'Semantic Addressing and Routing for Future Networks (SARNET-21) workshop report' (work-in-progress), IETF Information, Jun. 2021.

[Green2018] M. Green, 'Attack of the week: Group messaging in WhatsApp and Signal', Cryptography Engineering, Jan. 2018.

[Harris2021] J. Harris, 'Tim Berners-Lee: 'We need social networks where bad things happen less'', The Guardian, Mar. 2021.

[Hitz2021] S. Hitz, 'Next-generation Internet. Reliable. Secure.', Anapaya, May 2021.

[IETF/MLS] IETF/WG, 'Messaging layer security', accessed online in Jun. 2021.

[IETF-DetNet] N. Finn and P. Thubert, 'Deterministic Networking Problem Statement', IETF Information , May. 2019.

[IETF-http3] M. Bishop, 'Hypertext Transfer Protocol Version 3 (HTTP/3)' (work-in-progress), IETF Standard, Feb. 2021.

[IETF-IPv6-deployment] G. Fioccola and P. Volpato, 'IPv6 Transition Deployment Status' (work-in-progress), IETF Information, Oct. 2020.

[IETF-L4S] B. Briscoe, K. D. Schepper, M. Bagnulo and G. White, 'Low Latency, Low Loss, Scalable Throughput (L4S) Internet service: Architecture' (work-in-progress), IETF Information, Jul. 2021.

[IETF-MPQUIC] Q. D. Coninck and O. Bonaventure, 'Multipath Extension for QUIC (MP-QUIC)' (work-in-progress), IETF Standard, May 2021.

[IETF-RAW] P. Thubert, D. Cavalcanti, X. Vilajosana, C. Schmitt and J. Farkas, 'Reliable and available wireless technologies' (work-in-progress), IETF Information, Feb. 2021.

[IETF-RFC3272] D. Awduche, A. Chiu, A. Elwalid, I. Widjaja and X. Xiao, 'Overview and principles of Internet traffic engineering', IETF Information, May 2002.

[IETF-RFC4291] R. Hinden and S. Deering, 'IP version 6 addressing architecture', IETF Standard, Feb. 2006.

[IETF-RFC5681] M. Allman, V. Paxson and E. Blanton, 'TCP congestion control', IETF Standard, Sep. 2009.

[IETF-RFC7285] R. Alimi et al, 'Application-Layer Traffic Optimization (ALTO) protocol', IETF Standard, Sep. 2014.

[IETF-RFC7540] M. Belshe, R. Peon and M. Thomson, 'Hypertext Transfer Protocol Version 2 (HTTP/2)', IETF Standard, May 2015.

[IETF-RFC791] Information Sciences Institute University of Southern California, 'Internet protocol – DARPA Internet program – Protocol specification', IETF Standard, Sep. 1981.

[IETF-RFC793] Information Sciences Institute University of Southern California, 'Transmission control protocol – DARPA Internet program – Protocol specification', IETF Standard, Sep. 1981.

[IETF-RFC8312] I. Rhee, L. Xu, S. Ha, A. Zimmermann, L. Eggert and R. Scheffenegger, 'CUBIC for fast long-distance networks', IETF Information, Feb. 2018.

[IETF-RFC8446] E. Rescorla, 'The Transport Layer Security (TLS) protocol version 1.3', IETF Standard, Aug. 2018.

[IETF-RFC9000] J. Iyengar and M. Thomson, 'QUIC: A UDP-based multiplexed and secure transport', IETF Standard, May 2021.

[Inrupt] Inrupt, 'Solid – Sir Tim Berners-Lee's vision of a vibrant web, for all', accessed online in Jun. 2021.

[Kamiya2020] G. Kamiya, 'Data centres and data transmission networks', International Energy Agency Tracking Report, Jun. 2020.

[Katz2021] S. Katz, 'Google's Quic transmission protocol speeds up the Internet's flagship TCP', Tech Xplore, Jun. 2021.

[King&Farrel2021a] D. King and A. Farrel, 'Challenges for the Internet routing infrastructure introduced by changes in address semantics', IETF Informational, Jun. 2021.

[King&Farrel2021b] D. King and A. Farrel, 'A survey of semantic Internet routing techniques' (work-in-progress), IETF Information, Jun. 2021.

[Kumar2020] P. Kumar, 'QUIC - A Quick Study', arXiv, Oct. 2020.

[LiEtAl2020] Z. Li et al, 'Application-aware Networking (APN) framework' (work-in-progress), IETF Standard, May. 2021.

[Lohr2021] S. Lohr, 'He created the web. Now he's out to remake the digital world.', Chicago Tribune, Jan. 2021.

[Ofcom2021] The Office of Communications, 'Technology Futures – spotlight on the technologies shaping communications for the future', Jan. 2021.

[Rashid2020] F. Y. Rashid, 'Personal Virtual Networks [PVN] could give everyone more control over their data', IEEE Spectrum, Feb. 2020.

[SchinaziEtAl2020] D. Schinazi, F. Yang and I. Swett, 'Chrome is deploying HTTP/3 and IETF QUIC', Chromium Blog, Oct 2020.

[SolidMIT2017] Solid MIT, 'What is Solid?', Massachusetts Institute of Technology, 2017.

[SolidProject] Solid Project, 'About Solid', accessed online in Jun. 2021.

[SONiC] SONiC, 'What is SONiC?', accessed online in Jun. 2021.

[Treleaven2020] P. Treleaven et al, 'The pioneering distributed machine learning and privacy-preserving data technology'. UCL, 2020.

[W3C2015] W3C, 'Call for Participation in WebAssembly Community Group', W3C Community and Business Groups, Apr. 2015.

[W3C2020] W3C, 'W3C Solid community group', W3C Solid Community Group, Nov. 2020.

[W3Schools] W3 Schools, 'HTML Tutorial', accessed online in Jun. 2021.

[W3CSemanticWeb] W3C, 'The semantic web made easy', accessed online in Jun. 2021.

[WuEtAl2021] Q. Wu, Y. Yang, Y. Lee, D. Dhody, S. Randriamasy and L. Contreras, 'ALTO performance cost metrics', IETF Standard, Feb. 2021.

[ZigbeeC] Zigbee Alliance, 'Project Connected Home over IP [CHIP] connectivity standard developing into reality', accessed online in Jun. 2021.

R.6 Evolution in Internet compute, connect and cloud

[Ahmed&Taha2021] Z. K. Ahmed and F. H. Taha, 'Evaluation and realization of memristor emulator spice', Journal of Physics: Conference Series, 2021.

[ArmbrustEtAl2009] M. Armbrust et al, 'Above the clouds: A Berkeley view of cloud computing', Electrical Engineering and Computer Sciences University of California at Berkeley Technical Report, Feb. 2009.

[BytecodeAlliance] Bytecode Alliance, 'Mission [and] values', accessed online in Jun. 2021.

[CaleffiEtAl2018] M. Caleffi, A. S. Cacciapuoti and G. Bianchi, 'Quantum Internet: From communication to distributed computing', Proceedings of the ACM International Conference on Nanoscale Computing and Communication, Sep. 2018.

[Chopra2021] R. Chopra, 'Co-packaged optics and an open ecosystem', Cisco, Jan. 2021.

[Cloudflare] Cloudflare, 'What is BaaS? Backend-as-a-Service vs. serverless', accessed online in Jul. 2021.

[CloudHedge2020] Team CloudHedge, 'Container technology – A quick introduction, Jun. 2020.

[Djordjevic2012] I. Djordjevic, 'Quantum information processing and quantum error correction'. Academic Press, Apr. 2012.

[EA-Light2021] EurekAlert, ''Multiplying' light could be key to ultra-powerful optical computers', Feb. 2021.

[EA-Quantum2021] EurekAlert, 'Toshiba announces breakthrough in long distance quantum communication', Jun. 2021.

[Gartenberg2021] C. Gartenberg, 'IBM's first 2nm chip previews the processors of tomorrow', The Verge, May 2021.

[GoethalsEtAl2018] T. Goethals, M. Sebrechts, A. Atrey, B. Volckaert and F. De Turck, 'Unikernels vs containers: An in-depth benchmarking study in the context of microservice applications', IEEE International Symposium on Cloud and Service Computing, Nov. 2018.

[HamdiouiEtAl2015] S. Hamdioui, M. Taouil and H. A. D. Nguyen, 'Memristor: The enabler of computation-in-memory architecture for big-data', International Conference on Memristive Systems, Nov. 2015.

[Hennessy&Patterson2019] J. L. Hennessy and D. A. Patterson, 'A new golden age for computer architecture', Communications of the ACM, Feb. 2019.

[HPCwire2021] HPC wire, 'European Processor Initiative [EPI] announces EPAC1.0 RISC-V test chip taped-out', Jun. 2021.

[IET2021] Institution of Engineering and Technology, 'IBM unveils world's first 2nm chipmaking technology', May 2021.

[Intel] Intel, 'Neuromorphic computing', accessed online in Jun. 2021.

[Jain2020] M. Jain, 'Study of firecracker microVM', arXiv, May 2020.

[JoEtAl2010] S. H. Jo, T. Chang, I. Ebong, B. B. Bhadviya, P. Mazumder and W. Lu, 'Nanoscale memristor device as synapse in neuromorphic systems', Nano Letters, Mar. 2010.

[Johnson2021] D. Johnson, 'IBM introduces the world's first 2-nm node chip', IEEE Spectrum, May 2021.

[JonasEtAl2019] E. Jonas et al, 'Cloud programming simplified: A Berkeley view on serverless computing', Electrical Engineering and Computer Sciences University of California at Berkeley Technical Report, Feb. 2019.

[KingEtAl2003] S. T. King, G. W. Dunlap and P. M. Chen, 'Operating system support for virtual machines', Proceedings of the USENIX Technical Conference, Jun. 2003.

[LFEdge2020] The Linux Foundation Edge, 'Sharpening the edge: Overview of the LF Edge taxonomy and framework', LF Edge White Paper, Jul. 2020.

[Loughran2017] J. Loughran, 'Moore's law 2017: An uphill battle', The IET, May 2017.

[Maan&James2016] A. K. Maan and A. P. James, 'Voltage controlled memristor threshold logic gates', IEEE Asia Pacific Conference on Circuits and Systems, Oct. 2016.

[ManchesterU2018] The University of Manchester, ''Human brain' supercomputer with 1 million processors switched on for first time', News, Nov. 2018.

[Marena2018] T. Marena, '11 Myths about the RISC-V ISA', Electronic Design, Jan. 2018.

[Mell&Grance2011] P. Mell and T. Grance, 'The NIST definition of cloud computing', US National Institute of Standards and Technology Special Publication, Sep. 2011.

[MinkenbergEtAl2021] C. Minkenberg, R. Krishnaswamy, A. Zilkie and D. Nelson, 'Co-packaged datacenter optics: Opportunities and challenges', IET Optoelectron. Mar. 2021.

[Nayak2021] C. Nayak, 'Full stack ahead: Pioneering quantum hardware allows for controlling up to thousands of qubits at cryogenic temperatures', Microsoft, Jan. 2021.

[NN2020] Nature Nanotechnology, 'Beyond von Neumann', Jul. 2020.

[Ofcom2021] The Office of Communications, 'Technology Futures – spotlight on the technologies shaping communications for the future', Jan. 2021.

[OfcomSONIC2021] The Office of Communications, 'What is Open RAN, and why does it matter?', Jul. 2021.

[OsiaEtAl2020] S. A. Osia et al, 'A hybrid deep learning architecture for privacy-preserving mobile analytics', IEEE Internet of Things Journal, May 2020.

[Patrizio2021a] A. Patrizio, 'IBM leapfrogs everyone with its 2nm chips', Network World, May 2021.

[RISC-V] RISC-V International, 'About RISC-V', accessed online in Jun. 2021.

[SalimEtAl2019] S. S. Salim, A. Nisbet and M. Luján, 'Towards a WebAssembly standalone runtime on GraalVM', Proceedings Companion of the ACM SIGPLAN International Conference on Systems, Programming, Languages, and Applications, Oct. 2019.

[Shahsavari2013] M. Shahsavari, 'Memristor technology and applications: An overview', Microelectronic and Computer Engineering

Department Delft University of Technology Technical Report, Apr. 2013.

[Shepard2020] J. Shepard, 'Memory-centric computing and memory system architectures', Microcontrollers, Aug. 2020.

[Soni2018] D. Soni, 'Spiking neural networks, the next generation of machine learning', Towards Data Science, Jan. 2018.

[SONiC] SONiC, 'What is SONiC?', accessed online in Jun. 2021.

R.7 The Internet and energy efficiency

[Abbas&Yoon2015] Z. Abbas and W. Yoon, 'A survey on energy conserving mechanisms for the Internet of Things: Wireless networking aspects', Sensors, Sep. 2015.

[AjibolaEtAl2021] O. O. Ajibola, T. E. H. El-Gorashi and J. M. H. Elmirghani, 'Energy efficient placement of workloads in composable data center networks', Journal of Lightwave Technology, May 2021.

[Al-AzezEtAl2019] Z. T. Al-Azez, A. Q. Lawey, T. E. H. El-Gorashi and J. M. H. Elmirghani, 'Energy efficient IoT virtualization framework with peer to peer networking and processing', IEEE Access, Apr. 2019.

[Al-ShammariEtAl2019] H. Q. Al-Shammari, A. Q. Lawey, T. E. H. El-Gorashi and J. M. H. Elmirghani, 'Service embedding in IoT networks', IEEE Access, Dec. 2019.

[Al-QuzweeniEtAl2019] A. N. Al-Quzweeni, A. Q. Lawey, T. E. H. Elgorashi and J. M. H. Elmirghani, 'Optimized energy aware 5G network function virtualization', IEEE Access, Mar. 2019.

[AlharbiEtAl2020] H. A. Alharbi, T. E. H. Elgorashi and J. M. H. Elmirghani, 'Energy efficient virtual machines placement over cloud-fog network architecture', IEEE Access, May 2020.

[AliEtAl2017] H. M. M. Ali, T. E. H. El-Gorashi, A. Q. Lawey and J. M. H. Elmirghani, 'Future energy efficient data centers with disaggregated servers', Journal of Lightwave Technology, Oct. 2017.

[Allard2020] N. Allard, 'How will the new ITU standards affect the telecoms industry', Carbon Intelligence, Mar. 2020.

[Brown2019] C. Brown, 'Lithium ion batteries for the data center. Are they ready for production yet?', Uptime Institute, Nov. 2019.

[Burrows2021] L. Burrows, 'A long-lasting, stable solid-state lithium battery', Harvard School of Engineering and Applied Sciences, May 2021.

[Christian2020] A. Christian, 'Turns out dumping data centres in the ocean could be a good idea', Wired, Sep. 2020.

[DengEtAl2015] R. Deng, R. Lu, C. Lai and T. H. Luan, 'Towards power consumption-delay tradeoff by workload allocation in cloud-fog computing', IEEE International Conference on Communications, Jun. 2015.

[DongEtAl2011] X. Dong, T. El-Gorashi and J. M. H. Elmirghani, 'Green IP over WDM networks with data centres', Journal of Lightwave Technology, Apr. 2011.

[Dutta2020] A. Dutta, 'Security in SDN/NFV and 5G Networks', IEEE, Jan. 2020.

[ETSI2012] ETSI, 'Network functions virtualisation', ETSI White Paper, Oct. 2012.

[Faraday2020] The Faraday Institution, 'High-energy battery technologies', Jan. 2020.

[Gartenberg2021] C. Gartenberg, 'IBM's first 2nm chip previews the processors of tomorrow', The Verge, May 2021.

[GSMA&ITUEtal2020] Global System for Mobile Communications, 'Guidance for ICT companies setting science based targets', Science Based Targets Initiative, Feb. 2020.

[GSMA2021] Global System for Mobile Communications, 'Mobile net zero', Apr. 2021.

[Hawkins2019] A. J. Hawkins, 'Electric car battery with 600 miles of range? This startup claims to have done it', The Verge, Apr. 2019.

[HuEtAl2017] P. Hu, S. Dhelim, H. Ning and T. Qiu, 'Survey on fog computing', Journal of Network and Computer Applications, Nov. 2017.

[IsaEtAl2020] I. S. B. M. Isa, T. E. H. El-Gorashi, M. O. I. Musa and J. M. H. Elmirghani, 'Energy efficient fog-based healthcare monitoring infrastructure', IEEE Access, Oct. 2020.

[ITU2020] International Telecommunications Union, 'Greenhouse gas emissions trajectories for the information and communication technology sector compatible with the UNFCCC Paris Agreement', Jan. 2020.

[Judge2021] P. Judge, 'China launches underwater data center', Data Center Dynamics, Jan. 2021.

[LaweyEtAl2014a] A. Q. Lawey, T. E. H. El-Gorashi and J. M. H. Elmirghani, 'Distributed energy efficient clouds over core networks', Journal of Lightwave Technology, Jan. 2014.

[LaweyEtAl2014b] A. Q. Lawey, T. E. H. El-Gorashi and J. M. H. Elmirghani, 'BitTorrent content distribution in optical networks', Journal of Lightwave Technology, Aug. 2014.

[Lawrence2020] A. Lawrence, 'Data centre PUEs have been flat since 2013', Data Center Dynamics, Apr. 2020.

[Malmodin&Lundén2018] J. Malmodin and D. Lundén, 'The energy and carbon footprint of the global ICT and E&M sectors 2010–2015', Sustainability, Aug. 2018.

[MohamedEtAl2020] S. H. Mohamed, M. B. A. Halim, T. E. H. Elgorashi and J. M. H. Elmirghani, 'Fog-assisted caching employing solar renewable energy and energy storage devices for video on demand services', IEEE Access, Jun. 2020.

[MusaEtAl2018] M. O. I. Musa, T. E. H. El-Gorashi and J. M. H. Elmirghani, 'Bounds on GreenTouch GreenMeter network energy efficiency', Journal of Lightwave Technology, Dec. 2018.

[OsmanEtAl2014] N. I. Osman, T. El-Gorashi, L. Krug and J. M. H. Elmirghani, 'Energy-efficient future high-definition TV', Journal of Lightwave Technology, May 2014.

[Patrizio2021a] A. Patrizio, 'IBM leapfrogs everyone with its 2nm chips', Network World, May 2021.

[Patrizio2021b] A. Patrizio, 'The quest to run data centers on battery power', Data Center Dynamics, Apr. 2021.

[Poe] N. Poe, 'Lithium ion in the data center?', CoreSite, accessed online in Jun. 2021.

[Power2020] G. Power, 'The opportunity for telcos to act on climate', Macquarie, Apr. 2020.

[Roach2020] J. Roach, 'Microsoft finds underwater data centres are reliable, practical and use energy sustainably', Microsoft, Sep. 2020.

[Ye&Li2021] L. Ye and X. Li, 'A dynamic stability design strategy for lithium metal solid state batteries', Nature, May 2021.

[YosufEtAl2020] B. A. Yosuf, M. Musa, T. Elgorashi and J. Elmirghani, 'Energy efficient distributed processing for IoT', IEEE Access, Sep. 2020.

[YousefpourEtAl2019] A. Yousefpour et al, 'All one needs to know about fog computing and related edge computing paradigms: A complete survey', Journal of Systems Architecture, Sep. 2019.

Annex A: Acronyms

AI	Artificial intelligence
AIR	Automated Intelligent Regulation
AS	Autonomous Systems
BBR	Bottleneck bandwidth and round-trip
BGP	Border Gateway Protocol
CC	Creative Commons
CDN	Content distribution network
CPE	Customer premise equipment
CPU	Central Processing Units
CS	Conventional server
DC	Data centres
DetNet	Deterministic networking
EPI	European Processor Initiative
ETSI	European Telecommunication Standards Institute
FTP	File Transfer Protocol
FWA	fixed wireless access
GHG	Greenhouse gas
GIF	Graphics Interchange Format
GSMA	Global System for Mobile Communications
HTML	Hyper Text Markup Language
HTTP	Hypertext Transfer Protocol
HTTP/3	Hypertext Transfer Protocol Version 3
IBM	International Business Machine
ICT	Information and Communications Technology
IETF	Internet Engineering Task Force
L4S	Low Latency, Low Loss, Scalable
IoP	Internet of People
IoS	Internet of Senses
IoSr	Internet of Services
IoT	Internet of Things
ISA	Instruction Set Architecture
IP	Internet Protocol
IPv4	IP version 4
IPv6	IP version 6
ISD	Isolation Domain
ITU	International Telecommunication Union
JPEG	Joint Photographic Experts Group
LAN	local-area-network
LF	Linux Foundation
MLS	Message Layer Security
NFV	Network Function Virtualisation
NIN	Non-IP Networking
NOS	Network Operating System
OneDM	One Data Model
OSI	Open System Interconnect model
OSN	online social network
OS	Operating system
P2P	Peer-to-Peer
PS	Pod-scale
POD	Personal Online Datastore
PUE	Power usage effectiveness
PVN	Personal Virtual Network
QoE	Quality-of-Experience
QoS	Quality-of-Service
QUIC	Quick UDP Internet Connection
RAW	Reliable and available wireless
RISC-V	Reduced Instruction Set Computer
RPKI	Resource Public Key Infrastructure
RS	Rack-scale
SCION	Scalability, Control, and Isolation on Next-Generation Networks
Solid	Social Linked Data
SRv6	Segment routing over IPv6
SSI	Self-sovereign Identity
TCP	Transmission control protocol
TLS	Transport Layer Security
UDP	User Datagram Protocol
VM	Virtual machine
WAN	Wide-area-network

Annex B: Acknowledgements

We would like to offer our gratitude to the following experts who reviewed the whole or part(s) of the report:

- **Professor Jon Crowcroft,**
  Marconi Professor of Communications Systems, Cambridge University, UK.
- **Professor Dame Wendy Hall,**
  Regius Professor of Computer Science, University of Southampton, UK.
- **Dr. Kieron O'Hara,**
  Associate professor in Electronics and Computer Science, University of Southampton, UK.
- **Professor Jaafar Elmirghani,**
  Director of the Institute of Integrated Information Systems, University of Leeds, UK.
- **Dr. David Choffnes,**
  Associate Professor, Northeastern University, USA.

We also want to thank the subject matter experts listed below for their invaluable contribution and sharing of their expertise.

Roland Acra,
Chief Technology Officer,
Cisco, USA.

Sharad Agarwal,
Senior Principal Researcher,
Microsoft, USA.

Daniel Appelquist,
Director of Web Developer Advocacy,
Samsung, UK.

Jari Arkko,
Senior. Expert at Ericsson Research,
Ericsson, Finland.

Katja Bego,
Principal Researcher, Explorations & Futures,
Nesta, UK.

Benjamin Beurdouche,
Security Researcher,
Mozilla, USA.

Sir Tim Berners-Lee,
Director of the World Wide Web Consortium,
Inrupt, UK.

Scott Blue,
Spectrum Policy Expert,
Microsoft, USA.

Dr. Doug Burger,
Technical Fellow,
Microsoft, USA.

Brett Carr,
Manager DNS and Network Engineering,
Nominet, UK.

Vint Cerf,
Chairman, Marconi Society,
Google, USA.

Youssef Chami,
Head of Network Strategy,
Vodafone, UK.

Dr. David Choffnes,
Associate Professor,
Northeastern University, USA.

David Conrad,
Chief Technology Officer,
ICANN, USA.

Dave Coplin,
CEO and Founder,
The Envisioners, UK.

Professor Jon Crowcroft,
Marconi Professor of Communications Systems,
Cambridge University, UK.

Tom Curry,
Telco Cloud Architect
Google, USA.

Julie Dawson,
Director of Regulatory and Policy,
Yoti, UK.

Karl Dyson,
Enterprise Architect,
Nominet, UK.

Professor Jaafar Elmirghani,
Director of the Institute of Integrated Information Systems,
University of Leeds, UK.

Dr. Silvia Elaluf-Calderwood,
Consultant,
Strand Consult, USA.

Cath Goulding,
Chief Information Security Officer,
Nominet, UK.

John Grant,
Chairman of ISG Non-IP Networks and Next Generation Networks,
Nine Tiles, UK.

Dr. Hamed Haddadi,
Reader in Human-Centred Systems,
Imperial College, UK.

Professor Lynne Hall,
Professor of Computer Science,
University of Sunderland, UK.

Professor Dame Wendy Hall,
Regius Professor of Computer Science,
University of Southampton, UK.

Ted Hardie,
Vice President,
Cisco, USA.

Professor Natali Helberger,
Professor of Information Law,
University of Amsterdam, The Netherlands.

Marco Hogewoning,
Manager Public Policy and Internet Governance,
RIPE NCC, The Netherlands

Bobby Holley,
Distinguished Engineer,
Mozilla, USA.

Geoff Huston,
Chief Scientist**,**
APNIC, Australia.

Ankur Jain,
Distinguished Engineer and Senior Director,
Google, USA.

Tomas Jakimavicius,
Director, European Government Affairs**,**
Microsoft, USA.

Dr. Robert Joyce,
Chief Technology Officer,
Nokia Oceania, Australia.

Ari Keränen,
Expert in IoT Standards and Protocols,
Ericsson, Finland.

Mallory Knodel,
Chief Technology Officer,
Centre for Democracy and Technology, USA.

Konstantinos Komaitis,
Senior Director, Policy Development & Strategy,
Internet Society (ISOC), UK.

Pavan Kurapati,
Head of Global Service Provider Architects,
Juniper, USA.

David Lake,
Engineering Technologist,
Dell, UK.

Dr. Richard Li,
Chief Scientist and VP of Network Technologies,
FutureWei, USA.

Mike Liebhold,
Distinguished Fellow,
Institute for the Future, USA.

Julian Lucek,
Senior Distinguished Systems Engineer,
Juniper, USA.

Veronika McKillop,
President of the UK IPv6 Council,
UK IPv6 Council, UK.

Dr. Nicola Millard,
Principal Innovation Partner,
BT, UK.

Mark Nottingham,
Senior Principal Engineer
(Internet and Web Standards)
Fastly, Australia.

Dr. Kieron O'Hara,
Associate professor in Electronics and Computer Science,
University of Southampton, UK.

Dr. Colin Perkins,
Senior Lecturer (Associate Professor),
Glasgow University, UK.

Professor Dr. Adrian Perrig,
Swiss Federal Institute of Technology (ETH),
Switzerland.

Dr. Ioannis Psaras,
Researcher and Lecturer,
University College London, UK.

Professor Martin Reed,
University of Essex, UK.

Eric Rescorla
Chief Technology Officer, Firefox,
Mozilla, USA.

James Richards,
Security Researcher,
Nominet, UK.

Andrei Robachevsky,
Senior Director for Technology Programs,
Internet Society (ISOC), UK.

Peter Saint-Andre,
Senior Director, Strategic Partnerships,
Mozilla, USA.

Albert Sanders,
Global Head of Infrastructure Policy & Senior Counsel,
Google, USA.

Rowena Schoo,
Policy and Government Relations,
Nominet, UK.

Neal Smyth,
Head of DNS & Networks,
Nominet, UK.

Professor Melissa Terras,
Professor of Digital Cultural Heritage,
Edinburgh Futures Institute, UK.

Dave Thaler,
Software Architect,
Microsoft, USA.

Martin Thomson,
Distinguished Engineer,
Mozilla, USA.

Emmet Townsend,
Vice President,
Inrupt, UK.

Professor Philip Treleaven,
Director of the Financial Computing Centre,
University College London, UK.

Dr. Jo Twist,
Chief Executive Officer,
UKIE, UK.

Dr. Gareth Tyson,
Senior Lecturer,
Queen Mary University of London, UK.

Juan Carlos Zúñiga,
Head of Standardization & IPR Strategy,
Sigfox, France.

Annex C: Responses to our call for inputs

Below is the list of respondents who submitted non-confidential or partially confidential responses to our call for inputs. Their responses are available on the Ofcom website: https://www.ofcom.org.uk/consultations-and-statements/category-2/emerging-technologies.

Avelalto Ltd

BBC

BT

Enablersinc

Facebook

Federated Wireless Inc

GuRu Wireless

Internet Telephony Services Providers' Association (ITSPA)

JCDrawn

Lacuna Space

Lime Microsystems

Mastdata

Met Office

mmWave Coalition

N&M Consultancy Limited

Nextivity

SES

Starlink Internet Services UK

The Besen Group

The Telecom Infra Project

Viasat Inc