The Next Generation Internet

Anja Feldmann

Technische Universitat Berlin
Germany
anja@inet.tu-berlin.de

1. FUTURE INTERNET

Digitalization is transforming our world as we know it. This effects not only our society and science but also the economy. Indeed, digitalization is rewriting the rules of interaction and competition. While the digital technologies underlying these transformations are not necessarily new, in principle, they are going to be used at unthought-of scales and in unforeseen context.

In 10 to 20 years, staggering amounts of data will be available and include proprietary as well as public sources. Moreover, analytic and processing capabilities will have further advanced and offer intelligent machine learning mechanisms. In addition, we will take ubiquitous access to information from everywhere for granted.

We will not be bound to specific devices, rather, the most appropriate device that is in our immediate vicinity will be chosen based on our preferences. This allows us to interact with the information when and where ever we want. Such interactions will no longer require keyboards—they rely on speech, gestures, or emotion recognition, or even brain computer interfaces (BCI).

Thus, the mobile edge cloud will have become a reality. It is supported by micro data centers. Indeed, *virtualization* is everywhere, including network components, end-user devices, and eventually even sensors. Virtualization enables us to transparently use resources, including compute, storage, and network. However, the mobile edge cloud needs to be complemented by an appropriate world wide network as well as sufficient backend compute power and storage. Indeed, we will see a seamless integration of network, storage, and computing.

Driven by flexibility of which device to use when as well as the ease of infiltrating individual devices, as experienced with TeslaCrypt, we no longer store our private data on our own devices. One possible solution may be to use cryptographically secure storage chips. These may be coupled with cloud services for availability and/or backup purposes. This enables us to access data at every time from every where. Hence, the data is stored not only encrypted but also in a distributed manner. Moreover, the service may offer data availability "warranties" and possibly even an "insurance" against data loss and tampering.

The above is just one of the many services that will be enabled by the future Internet. To simplify application and service deployment even further, the next generation Internet will not only offer the simple "socket style" programming interface but also the concept of *CloudNets*. A CloudNet offers a single abstraction that provides an application with connectivity, storage, and processing capabilities. These resources are all bundled together and may provide guarantees. CloudNets provide isolation and their resource can be scaled out or down in a flexible manner as needed by the application. Thus, freeing the service operator from the details of the underlying physical infrastructure.

Hereby, all services have to operate on a *reliable* and *secure* infrastructure. This means, that not only each component by itself has to be secured and configured appropriately, but also the overall system does not have weaknesses. Thus, security has to make a significant step forward towards *usable security* in the sense that it can actually be used by everyone. Misconfiguration opportunities have to be minimized. Moreover, vigilant security awareness helps in minimizing problems and/or fix them as soon as they are noticed.

While putting an emphasis on reliability we also have to realize that problems and system failures cannot be avoided. Purely, due to scale individual components or even subsystems will fail. Thus, our toolbox will now include scalable methods that let us *debug* this complex Internet infrastructure and its services. We have to be able to trace service misbehavior to the responsible system component. We also need to determine which services are effected if a system component fails.

Moreover, we apply the same ideas and concepts to scalable data analysis. This allows us to trace the *data processing pipelines* and enables us to answer questions such as "based on which data elements was this information derived", "did anyone tamper with the data". Indeed, we will have many different interacting data processing pipelines which will have to be mapped to CloudNets and the appropriate network hardware. Hereby, we have novel elements that allow us to sample data, aggregate it, process it, as well as share results in a data privacy preserving manner.

We are currently at the beginning of the digital transformation in the sense that we have prototypical deployments of a subset of the above capabilities, e.g., Web search and advertisement. Yet, how to deploy and operate at scale is a huge challenge. This can be seen by the many open research opportunities in the context of Big Data, Data Pri-

vacy, Internet of Things/Internet of Everything/Industrial Internet/Tactile Internet. Additional challenges revolve around providing sufficient bandwidth (5G and beyond) and distilling information from data by intelligent data processing.

2. CURRENT INTERNET

While the Internet is currently viewed as widely successful for some of its participants, mostly the users and the content and service providers, e.g., Google, it still suffers from ossification in the underlying infrastructure. This ossification has multiple causes, among them is the fact that the Internet works quite well as it is. Therefore, Internet Service Providers (ISPs) have little incentive to change. After all, why should one change a running system? Moreover, ISPs suffer from a lack of business perspectives due to the predominant charging modi for Internet access: flat rates for users and a combined price model consisting of a base rate and usage based component for content providers. However, ISPs face substantial reductions of revenue combined with increased cost while content and service provides revenues have increased. This has started to motivate changes to the infrastructure as well as reorientation of some players. To enable future changes, e.g., to the infrastructure, we need *appropriate incentives* from the beginning.

Indeed, we see first changes in the infrastructure. IPv6 is finally becoming a reality since the IPv4 address space is exhausted. Moreover, we find that concepts such as software-defined networking and network-function virtualization are making their way into the infrastructure. In particular, mobile operators and service providers, e.g., Google, are moving along these lines. Within the next two deployment cycles—within 10 years—we can expect that the deployed hardware fully supports split forwarding architectures, enabled, e.g., by OpenFlow. These architectures separate control plane decision-making off from data plane forwarding. At that point the infrastructure is able to support custom programmability and partial centralization of the control plane, while allowing for commodity high-throughput, high-fanout data plane forwarding elements. However, to support the ever-growing hunger for more bandwidth we have to strive for a tighter integration of optical technology as well as inclusion of the access network.

With regards to security and misconfiguration opportunities we observe that protocol design and service development have come full circle. In the sense that initially the Internet was a cooperative environment. Once it was realized that the Internet was hostile it was presumed that it was possible to fence of services. However, attackers are getting stronger and assets more valuable and, thus, there is the attempt to get security "right". However, the drawback is complexity and, thus, we complete the circle back to a simple security model in a presumed friendly environment guaranteed by enhanced fencing mechanisms. However, fencing mechanisms do not suffice as highlighted by many security breaches in the past and present. Here, we finally need a distributed, yet simple and effective, security concept which encompases both the infrastructure and the services.

3. NEXT STEPS

Some of the above challenges stem from the original design of the Internet which was motivated by the need to connect supercomputer centers to enable easy data exchange. In such a world one does not have to worry about a multitude of different applications and data sets, misguided (wrong) information uncooperating users and/or security breaches, competition of providers, mobility, multitude of devices and sensors, information floods, closed user groups, or anonymous as well as authenticated communication. To enable the vision outlined above we need to further evolve the **Internet control plane**.

In addition, we have to develop scalable mechanisms to support **data processing pipelines**. Here, we need to tackle both: the **control** as well as the **data processing** plane.

Among the challenges that have to be addressed by both **control planes**—Internet and data processing pipeline—are the following:

- What are the right abstractions of the infrastructure for the control plane?
- What language do we use to describe the requirements and the infrastructure? How do we map between them?
- What functionality is hosted where inside the network?
- What service is mapped to which network function?
- How to specify what service guaranties are needed from the Infrastructure? How to map these to the needed resources? How to adapt the requirements over time and scale the service?
- How to automate the CloudNet deployment?
- How to enable debuggability and traceability?
- How to avoid/detect misconfigurations?
- What is the price/value of infrastructure components?
- Can we do revenue sharing between the service, content, and the infrastructure provider?

With regards to the data processing pipelines we also have to address questions of the **data processing plane**:

- How to find data that enables us to answer a question?
- How to combine data in intelligent manner to derive information via novel machine learning mechanisms?
- How and where to sample data/information?
- Where to store which data/information?
- Which access control mechanism should be used?
- How to keep the data from being misused?
- How to trace the data processing pipeline?
- How to provide long term scalable secure backup?
- How to enable information sharing?
- How do we price information? How do we price data? What are the resulting pricing models?

To realize ubiquitious access we have to tackle the questions of decentralized control. Due to the need for quick decisions not all decisions can be made centrally. Thus, we need control architectures which delegate some control decisions to decentralized controllers. This opens many tradeoffs and new interaction opportunities and threats. Indeed, we need this functionality across multiple operators and service providers. At the same time it needs to be invisible to the end-user.