Volume 21 • Issue 1 • November 2016

The Internet Touch

Fernando Kuipers

The vision of a tactile Internet includes transporting touch in addition to audio-visual interaction. In this article, we highlight our own work and some of the open challenges towards materializing that vision.

Touch me if you can

What if the Internet could be used to extend your touch in such a way that you would be able to interact with distant humans and cyber-physical systems[1] as if they were near? The potential applications would be endless and in some cases even might make the difference between life and death. One illustrative example is that of a surgeon who, via the Internet, controls a robotic arm to perform surgery on a patient in a different country. In general, transporting touch over the so-called Tactile Internet facilitates (1) transporting skills, for instance in health, design, manufacturing and repair, (2) education, through more intense and realistic learning experiences, and (3) enhanced social interaction, by literally being able to reach out. It may seem as science fiction, but devices already exist to transport touch, although presently they still lack the sensitivity we expect and are accustomed to.

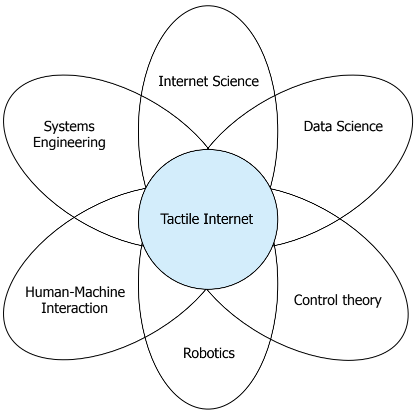

Figure 1: Many disciplines have to join forces to realize the Tactile Internet.

For all of the aforementioned application domains, the latency experienced should be very small, the connection should be highly reliable, haptic sensitivity should be present, etc. Indeed, as illustrated in Figure 1, realizing the Tactile Internet relies on consolidating advances in many disciplines. Robotics (or in this case the more appropriate term "cloud robotics") are needed to be able to control objects from a distance. As a result, also the field of human-machine interaction is applicable. There, the notion of "Quality of Experience" perceived by users needs to be revisited if touch is added to the equation. The fields of systems engineering, control theory, and data science are needed, since we are dealing with a system that is controlled by exchanging data. Our work, as we will describe in the following, fits under the umbrella of Internet Science.

[1] A cyber-physical system refers to an interacting network of physical and computational components, possibly controlled by humans from a distance.

The 1 ms challenge

In order to have real-time interaction, the round-trip latency between the communicating entities should not exceed 1 ms, else we perceive it as being unnatural. Unfortunately, even at the speed of light, this bounds the radius within which we would be able to communicate to 150 km. If we are to go beyond that boundary, this means we have to anticipate (and correct for) certain behavior, for instance via machine learning tools.

If a surgeon is operating remotely on a patient, a hick-up in the Internet connection could have devastating consequences. A highly reliable connection means that any failure of the connection cannot last longer than 1 ms per day. In that short time, the original connection needs to be restored or an alternative connection must have been set up. Note that this is in stark contrast to the 50 ms fail-over times currently permitted under "carrier-grade resilience."

A foundation for the future

In the current Internet, information from an application, like a web-browser, is cut into little pieces, which are put into IP packets, where IP stands for "Internet Protocol." These packets consist of a data part, i.e. the piece of information, and a header, which contains some fields, including one for the destination address. The packets are sent onto the Internet and are being forwarded by devices, such as switches and routers, based on the header. The switches and routers do their best to forward the packets to a next device closer to the destination, but generally cannot provide any guarantees on performance parameters like latency.

Since the Internet can be seen as an infrastructure for infrastructures, solving the challenges of the Tactile Internet will also provide a strong foundation for other, not necessarily touch related, domains that rely on the Internet. For example, smart buildings, smart energy grids, smart cities, e-health, autonomous vehicles, factories of the future, cloud robotics, etc.

Software-defined networking

Software-defined networking (SDN) is considered to be a revolution in how the Internet is controlled and may be just the technology we need to leap-frog towards the Tactile Internet. In SDN, see Figure 2, switches only perform data-plane functionality (that is, per-packet forwarding) and rely on a controller platform that provides the control-plane functionality (the "intelligence," like computing the forwarding rules and installing those in the switches). In contrast, traditional networking equipment embodies both functionalities into one physical device.

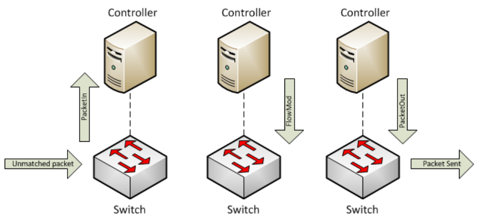

Figure 2: A switch (1) asks a controller if it knows how to handle a packet, (2) gets a rule installed by the controller if not, and (3) uses the rule to forward the packet towards the destination.

The decoupling of the data and control planes as in SDN, among other advantages, enables making tailor-made decisions on a per-flow[2] basis and offering guaranteed Quality of Service (QoS) – in contrast to the current best-effort service – for instance in terms of low latency.

The main potential of SDN lies in enabling modular network configuration through the use of network abstractions; just as operating systems for PCs rely on abstractions over multiple levels to manage their hardware resources. This modularity allows protocols etc. to be introduced/debugged/extended quickly and safely, instead of experiencing decade-long transitions (as for instance seen with IPv6).

In our work, we have developed OpenNetMon [1], a tool to measure QoS parameters in an SDN network, and we were the first to beat the 50 ms fail-over barrier needed for carrier-grade resilience [2], [3].

Stateful networking

The 1 ms requirements with respect to latency and reliability require us to anticipate certain events. Since, at a given time, one of multiple events may manifest, we have to prepare accordingly. As such, we have to provide different (SDN) rules for different states the network is in. For example, if a part of the network has been compromised for whatever reason, one may want to steer the network traffic in a different direction than for the uncompromised state. Hence, a prominent example where stateful networking would be useful is that of network resilience. Even at this moment, according to the Global Risks 2015 report from the World Economic Forum, critical information infrastructure breakdown is the most impactful technological risk. The report also states that the second most likely risk to occur (after interstate conflict) is that of extreme weather events, such as hurricanes. This, together with the even more stringent demands of the Tactile Internet, requires us to rethink how to protect our networks.

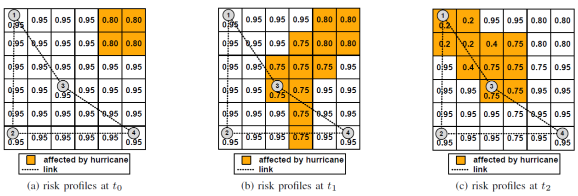

Figure 3: The network consists of four nodes in an area represented by a grid. The availability of the parts affected by the hurricane (in orange) reduces. In the last two time-slots, the availability of the network is affected.

As illustrated in Figure 3, in case of a disaster like a hurricane, a network that would be able to predict the hurricane's path and inflicted damage, might consider different network states over time, and as a result may mitigate large-scale connectivity disruption by steering its traffic over the most reliable connections. Our algorithms to compute those reliable paths are presented in [4]. By proactively changing paths, fail-over might be avoided altogether.

The concept of stateful networking does not only pertain to resilience, but can in principle link any event to a state. For instance, if a device is highly utilized, it might want to react differently than when it would have been idle. Or if a certain high-priority flow was just allocated, a device may want to sent other flows onto an alternative path to avoid interference. The network itself might even move to a different state by moving the locations of where certain network functions are executed, as in Network Functions Virtualization (NFV), where network functions (like firewall functionality) run in a virtualized environment on standard high-performance servers – akin to the principles of cloud computing – instead of on dedicated equipment per network function.

Like a puppet on a string

Wireless (5G) technology will be an important component of a Tactile Internet. After all, we want to maneuver freely and not be connected to a wire. However, in order to support the demand for higher and higher data rates, the wireless range over which we communicate must come down and, consequently, the number of these so-called smaller wireless cells will go up[3]. In our recent survey article [5], we describe how SDN has been used in the context of wireless networks and the Internet-of-Things and we list several open challenges.

One clear challenge relates to dealing with the increase in cells and devices. One proposal to address this scalability problem is called Fog Computing. There the compute, storage, and networking resources are not only available from dedicated servers, but also from user devices or at the edge of the network.

The road ahead

Clearly, there is more to say about the Tactile Internet than a 2-page article permits. We have simply provided a small sample of our own work that fits under the umbrella of the Tactile Internet. Recently, we (together with RangaRao Venkatesha Prasad) have initiated a joint project with the Indian Institute of Science, to collaborate and jointly supervise two PhD students on the topic of the Tactile Internet. If you would like to know more about the Tactile Internet, then you may want to start with the survey of Simsek et al. [6]. If you would like to actually work on the topic yourself, then feel free to contact me at F.A.Kuipers@tudelft.nl to discuss how we can join forces: after all, plenty of challenging problems still need to be solved...

References

[1] N. van Adrichem, C. Doerr, and F.A. Kuipers, "OpenNetMon: Network Monitoring in OpenFlow Software-Defined Networks," Proc. of the IEEE/IFIP Network Operations and Management Symposium, 2014.

[2] N. van Adrichem, B. van Asten, and F.A. Kuipers, "Fast Recovery in Software-Defined Networks," Proc. of the European Workshop on Software Defined Networking, 2014.

[3] N. van Adrichem, F. Iqbal, and F.A. Kuipers, "Backup rules in Software-Defined Networks," Proc. of the IEEE Conference on Network Function Virtualization and Software Defined Networks, 2016.

[4] F. Iqbal and F.A. Kuipers, "Spatiotemporal risk-averse routing," Proc. of the IEEE INFOCOM Workshop on Cross-Layer Cyber Physical Systems Security (CPSS), 2016.

[5] N. Bizanis and F.A. Kuipers, "SDN and virtualization solutions for the Internet of Things: A survey," IEEE Access, 2016.

[6] M. Simsek, A. Aijaz, M. Dohler, J. Sachs, and G. Fettweis, "5G-Enabled Tactile Internet," IEEE JSAC, 2016.

[2] A flow is a stream of packets belonging to the same application.

[3] Larger cells will likely co-exist, as they may serve different kinds of applications. For instance, LoRa is a novel technology for wireless sensor networks that specifically aims to connect many sensors at low data rates over long distances. If you want to know more about this, check out our new community at https://www.thethingsnetwork.org/community/delft/.