Accepted for publication in IEEE Vehicular Technology Magazine, to appear March 2014.

The Tactile Internet – Applications & Challenges

Gerhard P. Fettweis, Vodafone Chair Professor – TU Dresden, Germany

Wireless communications as known today enables to connect devices and people for exchanging content, being multimedia and/or data. The data rates of wireless communications continue to increase, mainly driven by innovation in electronics. An overlooked breakthrough is soon to happen, once the latency of communication systems becomes low enough to enable a round-trip delay from terminals through the network back to terminals of approximately 1ms. This is the response of human tactile to visual feedback control. In this case wireless communications can be the platform for enabling to control and steer real and virtual objects in many situations of our life. Almost no area of economy will be left untouched, with examples being health & care, mobility, education, manufacturing, smart grids, and many more. The Tactile Internet will become a driver for economies and innovation, and will help develop societies to a new level of sophistication.

Background and Introduction

Cellular communications as we know it today has been shaping the planet in an unprecedented way. Already the vast majority of people on the globe are connected by cellular communications. A currently happening next step of technology is to connect machines and devices as well. The success of cellular has flattened the planet and enabled economies to participate in the global economy, being the basis for many countries to become major players in world-wide trade and business. Also, cellular alone is still remaining to play a major part in impacting the economy of many countries, in particular in the developing world [1].

So far we have seen that cellular enables content to be moved around the globe, where prominent examples of content are voice telephony, text messaging, video streaming, emails, and files. The exponentially growing advances in electronics have led to an exponential increase in volume and size of the content to be transmitted. This again has resulted in the need for developing cellular standards that can handle continuously increasing data rates. Hence, to understand the future potential growth of wireless data rates, the future advances in electronics must be analyzed and understood.

However, is there another frontier to be tackled besides the race for data rate? When comparing the advances in orders of magnitude of increase in data rate with the reduction in round-trip latency of interaction, the latter has not dropped much below the requirement for telephony. LTE achieves a typical round-trip latency of 25ms [2], still far exceeding the 10ms requirement to start to enable real-time wireless gaming.

When moving to a round-trip latency of 1ms, and carrier grade robustness and availability, a new breakthrough in enabling unprecedented mobile applications becomes viable. These applications are coined as the "Tactile Internet", as this is the typical interaction latency required for a tactile steering and control of real and virtual objects, without creating cyber-sickness. This will revolutionize education, mobility & traffic, health & care, sports, entertainment, gaming, and the smart grid, just to name some segments which can be seen today already. The Tactile Internet will dramatically reshape our societies.

In the following at first a “sneak preview” on the future of electronics shall show us that we still have decades of fast innovations ahead. This will continue to drive data rates as well as computing capabilities along Moore’s Law. Then, the roadmap of wireless data rates will be drawn [3]. The latter is the basis to build the vision of the Tactile Internet. Extraordinary computing powers required for this will become available, based on the “sneak preview” given on electronics. Some application areas of the Tactile Internet are pointed out. Finally, a first list of key technical challenges ahead will be addressed, showing that wireless engineering remains to be a hot-bed for research.

A Sneak Preview on Future Electronics

Flash memories are the key storage technology for mobile devices of today. As the flash storage capacity is doubling every 18 months, this equals an increase by one order of magnitude (10x) every 5 years [4]. With the introduction of parallel machines, microprocessor processing capability has continued to experience the same exponential increase. Hence, as processing and memories continue to scale, so does the need for moving results/files, and therefore the crest for communications bandwidth is unrestrained.

The semiconductor technology scaling is slowing down, however, 3D integration is becoming available. “Yesterday” we designed silicon chips, wire or flip-chip bonded in a package. With the current advent of 3D chip integration as championed by Samsung [5] for memories so far, further advances will allow for 3D chip-stacks also for complex systems. This entails that the integration within a package to complex systems is bounded not by technology scaling but by the prospects of stacking many chips within a 3D chip-stack [6]. The end of integration is therefore not in sight, and the requirement for communications bandwidth not either.

However, one aspect is limiting this technology roadmap, which is the I/O bandwidth of devices. Chips or 3D chip-stacks need to be connected on a board, and boards via a backplane within a chassis. To address the interconnect challenges, solutions are being researched. Optical waveguides embedded in boards and optical integrated transceivers within chip-stacks can possibly enable breakthroughs in communications bandwidth. Furthermore, copper backplanes of today could be replaced by fully wireless chip-to-chip interconnectivity of boards in a chassis, again enabling increasing the bandwidth by more than three orders of magnitude beyond what we see today. This communications concept is behind building the "HAEC box" which is currently researched as the Highly Adaptive Energy Efficient Computing system [7].

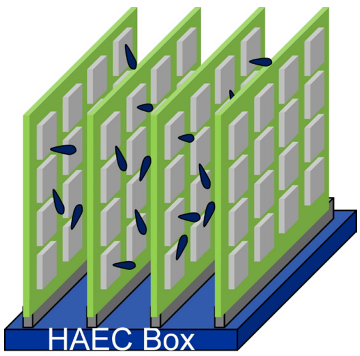

***Figure 1**: The vision of a HAEC box, "Highly Adaptive Energy Efficient Computing" box. – including embedded on-board optical interconnects and a wireless backplane by chip-stacks interconnecting with other chip-stacks on neighboring boards via a wireless (beam-steered) connection [7]. By e.g. having 4 boards as depicted with 16 chip-stacks within 10x10x10 cm^3, and every stack comprising of 128 chips with 128K cores each, 1 billion processing nodes could be packaged in a HAEC box.*

When combining 3D integration and possible advances in interconnect as sketched above, it could become feasible to build an exa-scale computing system with at least 1 billion compute nodes within a cube of one liter volume (10x10x10 cm^3). In Fig. 1 and [7] we refer to this as a "HAEC Box" vision. A consequence of this vision of a HAEC box is that an exa-scale amount of compute power could be placed in every access point or base station of future wireless systems. Hence, a future "Mobile Edge Cloud" system of unprecedented local compute power could become available, enabling many new applications and services that cannot be foreseen at this point in time.

Wireless Roadmap – The Race for Data Rate

Cellular Technology dominates today's life, and the appetite for bandwidth seems to be without foreseeable limits. Over the years we can observe an increase in data rate provided by wireless technologies by approximately 10x every five years, which is equivalent to Moore's Law of doubling every 18 months. As mentioned above, a key driver in need for capacity is most likely the continuous increase in storage and computing within mobile devices. And with the vision of a HAEC box possibly becoming true, advances in electronics remain to continue for more than the next two decades at a typical rate of Moore's Law. Hence, the appetite for more data rate over wireless will continue as well.

The introduction of each previous digital standard of the 3GPP family has been happening approximately a decade apart in time, driving the unrestrained race for data rate along this exponential increase.

Today it is widely accepted that LTE is a new generation of cellular communications, hence denoted as the 4th generation "4G". In its upcoming next intermediate step, referred to as LTE-Advanced, far more than 100 Mb/s data rate shall be achievable. When projecting a further unrestrained need for increase in data rate, and a 5th generation cellular system to become available in the decade of 2020-2030, cellular data rates of 1-10 Gb/s must become reality, see Fig.2.

Wireless LAN (local area network) technology has been driving along the same increase over the years, with a steady 10x improvement every 5 years. This requires a 100x higher data rate projection than cellular.

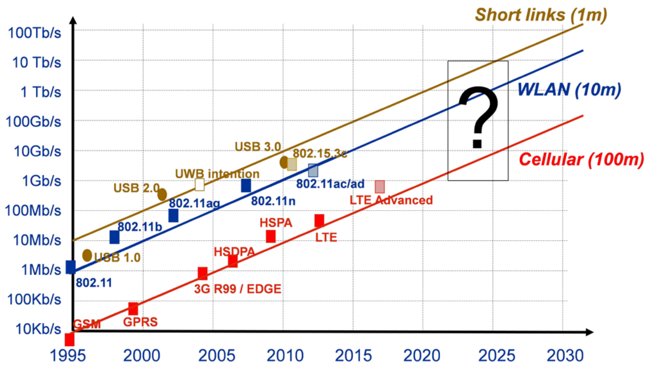

***Figure 2:** The Wireless Roadmap: Race for Data Rate. Here depicted are examples of standards [8].*

Clearly, we have seen an unprecedented line of applications becoming possible due to the available bandwidth. Examples which were unforeseeable 20 years ago are the mobile internet via smart phones, social networks, and search engines. However, data rate alone is not the only driver of innovation. When remembering the impact of the iPhone on the cellular market, it was not a leader in high bandwidth of connectivity, but was introduced with a cellular modem lagging a generation behind. It, however, had an unprecedented haptic interface, enabled by gyroscopes and a new quality level of its touch screen. Therefore the question must be posed if there is another major step ahead that could heavily change the way we interact? Within this paper the proposal is made that the Tactile Internet could play that kind of role.

The Tactile Internet – Real-Time Context

Real-time interaction with our environment is crucial to human beings. As humans we have many levels of defining real-time, since driving cars and making phone calls, and tactile touching of surfaces actuate many different sensors in our body system. To simplify, a brief review of a coarse granularity of different levels of real-time shall be discussed here. We define an interaction (service) to be real-time when the communication response time is faster than the time constants of the application, hence making the delay incurred by communication and computing negligible within the context considered. Here we only touch four types of physiological real-time constants, with further details given in [3,11]: muscular, audio, visual, and tactile. For literature on human reaction analysis we refer to [9], and the overview given in [10].

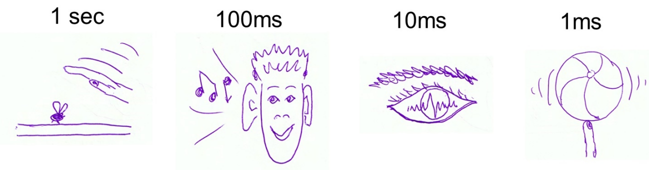

Figure 3: *A coarse grain analysis of human reaction and interaction times [3].*

How fast is fast? This question of the needs of round-trip latency must be answered within the context of a desired service and/or application, see Fig. 3. In case of muscular reaction times, as e.g. reacting on a fly settling on one of our limbs, the typical time is in the order of a second. When listening to voice and reacting within a conversation, latency of up to 100ms is non-observable. When watching a movie, a frame rate of 100Hz is acceptable, i.e. this relates to a latency of 10 ms between frames.

However, when tactile steering and/or control of an object is carried out and simultaneously seeing or hearing its reaction, a reaction latency of far below 10ms needs to be achieved. A simple example explaining this is given by moving an object on a touch screen. As a typical speed of our finger moving on a touch screen is 1m/s, a reaction time of the screen must be at 1ms to achieve an unnoticeable displacement of 1mm between the finger and the object to be moved. Assuming a latency of 100ms was implemented, this would lead to an unacceptable displacement of 100mm, i.e. 4 inches. Even 10ms latency which leads to 10mm displacement is clearly noticeable.

Similar studies on short latency requirements are known also from professional audio systems providers. Due to the fact that smaller than 5ms latency is difficult to achieve with digital communications, professional audio systems (e.g. for wireless stage microphones and mixing desks) are in majority analog. Another area in which this short latency constraint has been noticed is virtual reality and augmented reality. Synchronization between the ear and the eye has to be well below 10ms to prevent “simulator sickness”, which is now typically referred to as cyber sickness [12].

In case the feedback is given physically by a haptic system, it is referred to as “haptic communications” [13]. Here we refer to the Tactile Internet as systems which have tactile input but audio and/or visual feedback, requiring the extreme short latency.

The Tactile Internet – Application Scenarios

The scenarios above sketch a reasoning why a short 1ms reaction time is a new level of system performance. What is clear is that a *round-trip* latency is discussed, i.e. a tactile sensor reads information, and a connected system reacts with actuators seen by the same human within 1ms. This is why we refer to systems with this kind of latency as the “Tactile Internet”.

As current WLAN and cellular systems do not yield anything close to achieving a round-trip latency of 1ms, it therefore is difficult to comprehend all possible new applications which can emerge. In the following, some examples shall briefly give an idea on application scenarios, clearly showing the ground breaking potential of the new Tactile Internet.

Health & Care:

The UC Berkeley’s Robotics and Human Engineering Laboratory has demonstrated the idea: people who were tied to using a wheel chair are able to walk with exoskeletons. This has drawn a lot of attention world-wide, now also showing first results in making patients with cross section palsy being able to walk again [14]. Exoskeletons are devices which are strapped to your body and limbs, enabling to empower you to move your limbs with more force than available from your muscles. The reaction time of an exoskeleton has to be such that movements are within tactile latency. We can envision that no elderly person has to use a walker anymore, but instead strap on exoskeletons. These must be controlled by wireless systems recognizing other people with/without exoskeletons as well as the surroundings. In both cases this is for not harming others or the surroundings, and e.g. not falling down stairs and running into objects with the force of an exoskeleton. However, if we can build wireless systems with tactile latency, a whole new world can open up for physically impaired people. Independent of your current physical ability, many people can then freely move around again.

Besides this, many applications in the manufacturing and construction industry can be foreseen by the use of exoskeletons. Again, in many cases a tactile latency (clearly below 10 ms) of interaction with other objects must be enabled to make the application safe.

Education & Sports

This concept of using exoskeletons connected wirelessly can also be used for the training of physical movements. Possibilities could be for learning movements for sports (e.g. surfing), or for maneuvering vehicles or operating machines in “virtual training centers” before going live. In addition this can be used for remote physiotherapy, i.e. strapping on exoskeletons and being connected to a therapist, executing training sessions without having to physically visit the therapist’s office.

The possible variety of applications is very large for this educational concept enabled by the Tactile Internet. It will drive education to a new level of possibilities. Examples are interactive virtual history sessions, where a whole class room jointly engages via a virtual reality setting in experiencing the old Rome.

Traffic

The Tactile Internet with approximately a 1ms round-trip delay could revolutionize mobility. Assume that in cities all vehicles were driving at 10m/s (36km/h or 20mph), then a car could come to full-stop within 10m, i.e. 2 seconds. Fully automatic driving would be possible, making traffic lights superfluous, as vehicles could zip through intersections without crashing. Also, pedestrians could engage a personal bubble with their cell phone which makes sure that no car hits the person when crossing roads.

Platooning of vehicles is the joint formation of a line of cars travelling at the same speed with a very short distance from each other. Its introduction could generate great fuel efficiency, but this requires tight stability control of the platoon of cars. Today the electronic stability control (ESC) also referred to as electronic stability program (ESP) of a vehicle must be synchronized between wheels within 1-2ms, to ensure that the car does not lose stability. However, in the case of a platoon of cars with ESC/ESP, the complete platoon needs to have coordinated action of the wheels.

Hence, all vehicles of a platoon need to be connected with 1-2ms latency.

Another application could be remote driving. If the weather makes the driver feel unsafe, a call center shall be available for taking over the controls of the car and remotely driving it to the desired destination. However, tactile control must be possible to ensure real-time feedback for safe remote driving.

Robotics and Manufacturing

Controlling robots must happen at latency reaction times that are fast enough for the robot and its object not to start mechanically resonating. For many robotics scenarios in manufacturing this has led to a maximum latency target of a communication link of 100µs, and round-trip reaction times of 1ms, the target as discussed for the Tactile Internet.

Henry Ford introduced the assembly line for mass production of one model. Today, a diversified product portfolio should be produced within one assembly environment. This, however, can make robots remain idle at times when a product being assembled does not require its engagement. Hence, a future factory could have mobile robots join assembly lines or assembly points just in time for their action. In this case, a wireless infrastructure of a Tactile Internet will become necessary.[1]

Free-Viewpoint Video [15]

Digital image processing of video information can be used to synthetically render the viewpoint of the viewer to another spot. In future the choice of point-of-view can be dynamically modified, allowing the projection of an extended view beyond one's current viewpoint. An example of interest could be viewing a sports event in a stadium, where the stadium is equipped with hundreds of cameras. Real-time rendering of these cameras allows each person to choose their viewpoint of interest, e.g. the viewpoint of the favorite player on the field. This is projected onto the personal tablet PC or smart phone carried along by the stadium visitor. To maximize the excitement, the latency between the naturally visible surroundings by a visitor and the free viewpoint video being displayed on the tablet must be minimized. Again a 1ms latency would generate a fully synchronized experience without creating cyber sickness.

Also, video cameras of cars and smart phones could be connected and rendered into a free viewpoint video, not only in stadiums, but between cars on a road, or between vehicles in logistics hubs, or between pedestrians in crowded areas. This would always allow for projecting one's own viewpoint into a free viewpoint of choice at tactile latency.

Smart Grid

Another application area is synchronization of suppliers in a smart grid. As synchronous co-phasing of suppliers is necessary to minimize reactive power, this must be achieved within a small angle of phase. A 1ms latency of communications between (local) suppliers can ensure an 18° (50Hz AC network) or 21.6° (60Hz) phase coherence.

[1] For industrial manufacturing communications requirements see http://www.bmbf.de/foerderungen/22967.php

The Paradigm Shift

The Tactile Internet is a big paradigm shift lying ahead. It enables us to engage with objects of our environment at tactile latencies. Today communications technology is widely used for moving content from A to B. Content can be multimedia content or static content, e.g. video streams, data files, voice, or email. By making the Tactile Internet happen, we can engage and steer real and virtual objects directly at tactile real-time interaction speeds. Hence, in this case the primary goal of communications is to provide the necessary infrastructure for this. Not content needs to be transported, but control information. The paradigm shift of the Tactile Internet is therefore that communications is built for enabling steering and control, a big differentiator from moving content as today.

The Fig. 4 depicts this paradigm shift from content to controls. This opens up completely new opportunities for existing and new applications in many fields, of which the ones listed in Fig. 4 are only a few.

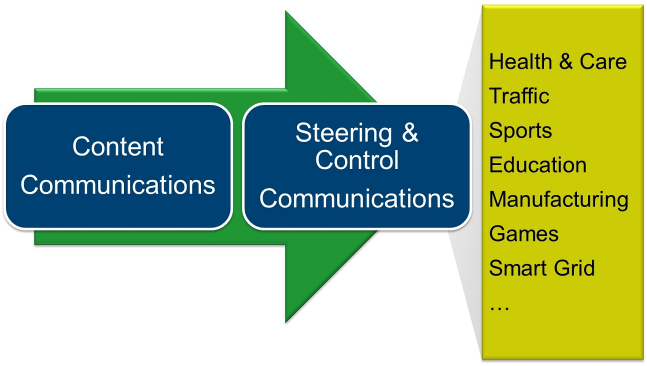

***Figure 4:** The road ahead enabled by the Tactile Internet – stepping from moving content to controlling and steering real and virtual objects of our environment. This enables applications in many segments, of which a few are listed.*

Some Key Challenges Ahead

Obviously, requiring a 1ms round-trip latency is a huge challenge. The physical transmission must have very small packets, to enable e.g. a one-way physical layer transmission of 100µs. To achieve this, each packet cannot exceed 33µs packet duration. The reason for this being that packet error correction encoding at the transmitter, and error correction decoding and detection at the receiver limit the packet size to less than 1/3 of the target latency. This clearly indicates that the modulation used in current LTE cellular systems is not a viable proposal, as the duration of one OFDM symbol alone is on the order of 70µs long. A complete rehaul of the cellular physical layer is necessary for the Tactile Internet, which might become reality with the 5th generation system.

Also, current systems have little provisioning for defining the minimum latency of protocol handling. A TCP/IP packet arriving at a base station which is labeled to tactile latency must be routed directly to the physical layer for transmission. Today, the latency of wireless communications is mainly analyzed by evaluating the structure of the medium access control (MAC) as well as by the latency required for the signal processing of the physical layer. However, protocol processing will become a major challenge for making the Tactile Internet possible. Steering and control packets arriving at a base station must be handled immediately and "tunneled" into reserved MAC slot. It certainly looks very much more like a circuit switching challenge in terms of latency and packet reservation guarantees than today's mainstream packet switching paradigm.

Another key challenge is to be able to provide carrier grade access reliability and robustness. Fixed-line carrier grade requires an uptime of the system of so-called seven nines, being 99.99999%, i.e. a failure rate of 10^{-7}. If we move to having objects of our environment being steered and controlled by the Tactile Internet, reliability of a link as well as availability of setting-up connectivity are of utmost importance. A failure rate even below 10^{-7} might be necessary in some cases. How can this small number of merely 3.17 seconds a year outage be achievable? Wireless systems are today built under the impression that a link with 3% outage is a good link. This is far from 10^{-7}. However, when two links with uncorrelated channels are combined, 3% outage per link generates a combined outage of approximately 10^{-3}. And five uncorrelated links can already achieve carrier grade outage of less than 10^{-7}!

Today a hot topic of cellular systems is carrier aggregation. In case of inter-band carrier aggregation it can be assumed that the links have uncorrelated channels. Another hot topic in cellular system design of today is CoMP (coordinated multipoint [16]) or ICIC (interference cancellation and interference coordination). In the special case of achieving this by distributed MIMO the mobile must be connected to at least two base stations simultaneously. Again this creates parallel uncorrelated links, which are used today for increasing the data rate and fairness. However, these uncorrelated links can also be aggregated for reducing the outage instead. Hence, building wireless connectivity that achieves carrier grade outage must not be an insurmountable problem, but can be realized by using access diversity either due to connecting to multiple base stations or via multiple different carriers simultaneously.

To achieve a round-trip latency of 1ms, the communication delay due to the speed of light needs to be considered as well. Within 1ms light travels 300km, i.e. the maximum distance for a steering and control server to be placed from the point of tactile interaction by the users is 150km away. However, this assumes no processing delays anywhere along the communication. Taking the additional signal processing, protocol handling, and switching delays into account, this requires the control/steering server to be in the range of 15km from the tactile point of interaction.

Extremely powerful servers, as e.g. the HAEC box, therefore have to be built as close to the base station as possible, i.e. at the edge of the mobile radio access network. This has been coined as the "mobile edge cloud" concept. The possibly best solution is to combine servers into the same box as base stations and access points. In addition, HAEC box like servers must have real-time operating systems which can guarantee an extremely small response time. The highly parallel server system must also

Users can be mobile, requiring a handoff of the communication link(s) from one base station to another. However, an application connected to the user will be running on a local HAEC box like server. Hence, in addition to the communications handoff, the hardware and software system running on these servers must also be built for a handoff of a running application from one server at one base station location to the next one.

Finally, a challenge which should not be left without mentioning is providing safety and security of the Tactile Internet and its applications. Clearly, if we rely on the Tactile Internet providing daily applications which can cause great harm if not operating properly, the Tactile Internet must be very secure and safe. Also, for many reasons, data integrity and accountability must be provided as well.

Conclusions

Today's wireless communication is designed for transporting content. The breakthrough lying before us is to enable wireless steering of our environment, enabling the Tactile Internet[2]. For this some key features must be achieved, of which most notably a short round-trip latency of approximately 1ms is necessary. In addition 99.99999% carrier grade robustness and availability must be provided, unheard of in today's systems. Once the challenges have been addressed and solutions have been found, the Tactile Internet will enable an unforeseeable plurality of new applications, products, and services. Discovering this world of the Tactile Internet will become a main driver for innovation and for our economies. Societies will change, with many new services becoming available which will make the planet a better place.

References

[1] Vodafone Institut, "Mobile Technologies – the digital fabric of our lives," Sep/2013, www.vodafone-institut.de/uploads/media/Studies_komplett_FINAL_von_250913.pdf

[2] Ericsson white paper, "LTE for Utilities – Supporting Smart Grids," Uen 285 23-3208, September 2013.

[3] G.P. Fettweis, "A 5G Wireless Communications Vision," Microwave Journal, December 2012, pp. 24-36.

[4] "International Technology Roadmap for Semiconductors," 2012 Edition, www.itrs.net/Links/2012ITRS/Home2012.htm

[5] Chang-Gyu Hwang, Keynote, IEEE IEDM, San Francisco, 2006. www.samsung.com/us/news/3435.

[6] R. Puschmann, "3D Integration of standard integrated circuits," IEEE 3DIC Conference, San Francisco, USA, Oct. 2-4 2013.

[7] G.P. Fettweis, W. Lehner, W. Nagel, "Pathways to Servers of the Future: Highly Adaptive Energy Efficient Computing (HAEC)," in Proc. of DATE-Conference 2013, pp. 1161-1166.

[8] G.P. Fettweis, "LTE: The Move to Global Cellular Broadband," Foreword, to appear in Intel Technical Journal, special issue on LTE, 2014.

[9] Jeff Johnson, "Designing With the Mind in Mind," Morgan Kaufman Publ., 2010.

[10] T. DeFanti and R. Stevens, "Teleimmersion," Chapter 6 within The Grid: Blueprint for a New Computing Infrastructure, The Elsevier Series in Grid Computing, pp 131-155.

[11] G.P. Fettweis and S. Alamouti, "5G – Personal Mobile Internet Beyond What Cellular Did to Telephony," accepted for IEEE Communications Magazine, to appear in February 2014.

[12] G.C. Burdea and P. Coiffet, "Virtual Reality Technology," IEEE Press, 2003.

[13] E. Steinbach, S. Hirche, M. Ernst, F. Brandi, R. Chaudhari, J. Kammerl, I. Vittorias, "Haptic Communications," Proc. of IEEE, Vol. 100, No. 4, April 2012, pp. 937-956.

[14] Die Auferstehung, Die Zeit, www.zeit.de/zeit-wissen/2013/02/Exoskelette-Querschnittslaehmung, March-7, 2013.

[15] J. Carranza et al., "Free-viewpoint video of human actors," ACM Transactions on Graphics (TOG) - Proceedings of ACM SIGGRAPH 2003, Volume 22 Issue 3, July 2003 pp. 569-577.

[16] P. Marsch and G.P. Fettweis, "Coordinated Multi-Point in Mobile Communications: From Theory to Practice," Cambridge University Press, August 2011.

2 G.P. Fettweis, "The Tactile Internet – Enabled By 5G," Youtube, www.youtube.com/watch?v=VXEPzQgpok, keynote at 2013 Johannesberg Summit.

Acknowledgements

The author would like to thank many co-researchers who served as sounding board and sharpened many ideas put forward in this paper. Special recognition shall be made towards Siavash Alamouti, Thomas Wiegand, Steffen Watzek, Erich Zielinski, Holger Boche, my whole team at TU Dresden, and the Vodafone Group R&D team.

Gerhard P. Fettweis earned his Ph.D. under H. Meyr at RWTH Aachen in 1990. Thereafter he was at IBM Research and then at TCSI Inc., California. Since 1994 he is Vodafone Chair Professor at TU Dresden, Germany, with main research interest on wireless transmission and chip design. He is IEEE Fellow, and honorary doctorate of TU Tampere. As repeat entrepreneur he has co-founded 11 startups so far.
He setup funded projects of close to EUR 1/2 billion volume, notably the German science foundation's CRC HAEC and CoE CfAED. He helped organizing IEEE conferences, e.g. as TPC Chair of ICC 2009, General Chair of VTC Spring 2013 and DATE 2014.