



The Tactile Internet

ITU-T Technology Watch Report
August 2014

Extremely low latency in combination with high availability, reliability and security will define the character of the *Tactile Internet*. It will have a marked impact on business and society, introducing numerous new opportunities for emerging technology markets and the delivery of essential public services. This Technology Watch report outlines the potential of the *Tactile Internet*, exploring its promise in application fields ranging from industry automation and transport systems to healthcare, education and gaming. It goes on to describe the *Tactical Internet's* demands on future digital infrastructure and its expected impact on society, concluding with a brief discussion of the role to be played by the ITU framework.



The rapid evolution of the telecommunication/information and communication technology (ICT) environment requires related technology foresight and immediate action in order to propose ITU-T standardization activities as early as possible.

ITU-T Technology Watch surveys the ICT landscape to capture new topics for standardization activities. Technology Watch Reports assess new technologies with regard to existing standards inside and outside ITU-T and their likely impact on future standardization.

Acknowledgements

This report was written by:

Prof. Gerhard Fettweis – Technische Universität Dresden
Prof. Holger Boche – Technische Universität München
Prof. Thomas Wiegand – Technische Universität Berlin and Fraunhofer Heinrich-Hertz-Institute
Prof. Erich Zielinski – Alcatel-Lucent Stiftung für Kommunikationsforschung
Prof. Hans Schotten – DFKI and University of Kaiserslautern
Peter Merz – Nokia Solutions and Networks Management International GmbH
Prof. Sandra Hirche – Technische Universität München
Dr. Andreas Festag – Technische Universität Dresden
Dr. Walter Häffner – Vodafone GmbH
Dr. Michael Meyer – Ericsson GmbH
Prof. Eckehard Steinbach – Technische Universität München
Prof. Rolf Kraemer – IHP, Innovations for High Performance Microelectronics
Prof. Ralf Steinmetz – Technical University Darmstadt
Dr. Frank Hofmann – Robert Bosch GmbH
Prof. Peter Eisert – Fraunhofer Heinrich-Hertz-Institute
Dr. Reinhard Scholl – International Telecommunication Union
Prof. Frank Ellinger – Technische Universität Dresden
Dr. Erik Weiß – Telekom Deutschland GmbH
Ines Riedel – Technische Universität Dresden

The authors highly appreciate the support from the *Alcatel-Lucent Stiftung für Kommunikationsforschung* and from the project “*fast strategy*” as part of the cluster “*fast – Fast Actuators Sensors and Transceivers*” in the framework of the initiative “*Zwanzig20 – Partnerschaft für Innovation*” of the Federal Ministry of Research and Education (BMBF).



Please send your feedback and comments to tsbtechwatch@itu.int.

The opinions expressed in this report are those of the authors and do not necessarily reflect the views of the International Telecommunication Union or its membership.

© ITU 2014

All rights reserved. No part of this publication may be reproduced, by any means whatsoever, without the prior written permission of ITU.

Table of contents

	<i>Page</i>
Executive Summary	1
Introduction	2
1. Vision <i>Tactile Internet</i> – An Important Future Innovation Area	2
2. Impact on Society	4
3. Application Fields	5
3.1 Industry	5
3.2 Robotics and Telepresence	5
3.3 Virtual Reality	6
3.4 Augmented Reality	7
3.5 Healthcare	8
3.6 Road Traffic	8
3.7 Serious Gaming	9
3.8 Education and Culture	10
3.9 Smart Grid	11
3.10 Additional Application Fields	11
4. Infrastructure Requirements	12
4.1 Latency and Reliability	12
4.2 Security	13
4.3 System Architecture	14
4.4 Sensors and Actuators	15
4.5 Access Networks	16
4.6 Mobile Edge-Cloud	16
5. ITU Framework	17

Executive Summary

Part 1 and 2: The vision of the Tactile Internet and its impact on society

The Tactile Internet will find novel application fields in which to contribute to the solution of the complex challenges faced by our society, with a sample of its value to social wellbeing found in education and lifelong learning, healthcare, personal safety zones, traffic in a smart city or energy.

The Internet of Things connects devices, or objects, to increase their efficiency by exploiting the potential of networking. The next wave of innovation will create the Tactile Internet, introducing numerous new opportunities for emerging technology markets and the delivery of essential public services.

In principle, all of our human senses can interact with machines, and technology's potential in this respect is growing. The Tactile Internet will enable haptic interaction with visual feedback, with technical systems supporting not just audiovisual interaction, but also that involving robotic systems to be controlled with an imperceptible time-lag.

Part 3: Application Fields - The potential of the Tactile Internet

Emerging Tactile Internet application fields are progressing towards precise human-to-machine and machine-to-machine interaction, with key examples found in industry, robotics and telepresence, virtual reality, augmented reality, healthcare, road traffic, serious gaming, education and culture, and smart grid.

1-millisecond end-to-end latency is necessary in Tactile Internet applications. For technical systems to match humans' interaction with their environment, our natural reaction times set the targets that technical specifications must meet.

The high availability and security, ultra-fast reaction times and carrier-grade reliability of the Tactile Internet will add a new dimension to human-to-machine interaction by delivering a latency low enough to build real-time interactive systems. The professional digital infrastructure will similarly revolutionize machine-to-machine interaction.

Part 4 and 5: Infrastructure requirements & ITU framework

The technical requirements of the Tactile Internet place extraordinary demands on future networks' latency and reliability, security, system architecture, sensors and actuators, access networks and mobile edge-clouds.

Existing infrastructures are both technically and conceptually insufficient to support the envisioned applications of the Tactile Internet, with certain areas in need of accelerated research and development. In parallel, supporting standardization systems are expected to evolve in line with the diversity of the Tactile Internet's application areas and associated stakeholder ecosystem.

Given ITU's vital role in the global management of the wireless frequency spectrum and satellite orbits, its global reach including developing countries, and the wealth of experience in defining new-technology landscapes through the publication of the requisite standards, ITU's unique public-private membership model offers an ideal platform for collaboration open to all invested in realizing the vision of the *Tactile Internet*.

Introduction

Mobile data communication is omnipresent. The *Mobile Internet* connects people anywhere and allows for voice services and the exchange of data and multimedia content at any time. Numerous innovations in the information and communication technology (ICT) sector have enabled exponential growth in network capacity, leading to the emergence of smartphones and a user-experience rich in multimedia.

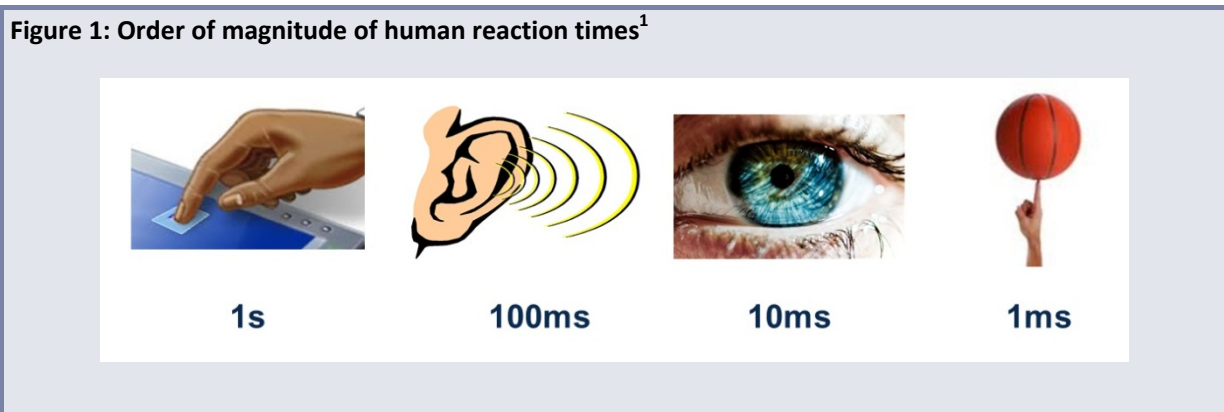
The *Internet of Things (IoT)* connects devices, or objects, to increase their efficiency by exploiting the potential of networking. The next wave of innovation will create the *Tactile Internet*.

Extremely low latency in combination with high availability, reliability and security will define the character of the *Tactile Internet*. It will have a marked impact on business and society, introducing numerous new opportunities for emerging technology markets and the delivery of essential public services.

This report outlines the potential of the *Tactile Internet*, exploring its promise in application fields ranging from industry automation and transport systems to healthcare, education and gaming. It goes on to describe the *Tactile Internet's* demands on future digital infrastructure and its expected impact on society, concluding with a brief discussion of the role to be played by the ITU framework.

1. Vision Tactile Internet – An Important Future Innovation Area

Human beings' interaction with their environment is crucial. Our senses allow us to perceive our environment and decide whether to adapt ourselves to that environment or modify it (changing our environment). Our perceptual processes limit the speed of our interaction with our environment. We experience interaction with a technical system as intuitive and natural only if the feedback of the system is adapted to our human reaction time. Consequently, the requirements for technical systems enabling real-time interactions depend on the participating human senses.



Human reaction times depend on the sensory stimulus and whether the human is prepared or unprepared for the situation (Figure 1).

When reacting to a sudden, unforeseen incident, the time-lag between a human sensing a stimulus and responding with a muscular reaction is in the range of 1 second.

For technical applications to match humans' interaction with their environment, our natural reaction times set the targets that technical specifications must meet.

¹ Image: Fettweis, G.; Alamouti, S., "5G: Personal Mobile Internet beyond What Cellular Did to Telephony," *Communications Magazine, IEEE*, vol. 52, no. 2, pp. 140-145, February 2014

An intuitive example is interactive web browsing. To experience immediacy, the page build-up after clicking on a link should be a fraction of the human unprepared reaction time. Real-time experience for browsing interaction is achieved only if a new web page can be built-up within a few hundred milliseconds of a user clicking on a hyperlink.

If a human is prepared for a situation, it is clear that a faster reaction time is needed.

The human auditory reaction time is about 100 milliseconds. To enable natural conversation, modern telephony is designed to ensure that voice is transmitted within 100 milliseconds. Higher latencies would disturb us.

A typical human visual reaction time is in the range of 10 milliseconds. To allow for a seamless video experience, modern TV sets have a minimum picture-refresh rate of 100 Hertz, translating into a maximum inter-picture latency of 10 milliseconds.

But if a human is expecting speed, such as when manually controlling a visual scene and issuing commands that anticipate rapid response, 1-millisecond reaction time is required. Examples are moving a mouse pointer over a screen and viewing a smooth path of the pointer over the screen, or moving our heads while wearing Virtual Reality (VR) goggles and expecting an immediate response from the visual display.

In principle, all of our human senses can interact with machines, and technology's potential in this respect is growing.

Among the senses, the visual-tactile interaction between humans and technical systems is becoming more important, especially in light of the proliferation of smartphones. Alongside the rapid data transmission, there are two other critical questions to be addressed in human-to-machine interaction: in order to be intuitive and natural, the reaction of the technical system must fit within the window framed by the reaction time of the considered human sense. When multiple senses are participating in the interaction, the time-lag of the feedback for different senses has to be imperceptible. Too much of a time-lag between the visual and auditory stimulus would disturb us, as in the case of a missed lip-synchronization.

The most challenging latency requirement for technical systems arises in tactile or haptic interaction – our sense of touch and the movement of our limbs interacting with visual or auditory feedback. Think of moving a 3D object with a joystick or in a virtual-reality environment. The introduction of the human body and movement in this equation results in strict latency requirements, in the order of 1 millisecond. If the time-lag between the virtual picture and human movement is above 1 millisecond, 'cybersickness' may occur², with users becoming disoriented in an experience similar to the motion sickness sometimes suffered at sea, in the air or on the road.

Technical systems have a so-called 'end-to-end latency', which comprises all delays experienced by a communication from origin to destination. It includes the time spent in the transmission of the information from a human via the communication infrastructure to a control server; the processing of the information and generation of a reaction; and the eventual transmission of the reaction via the communication infrastructure back to a human. If the end-to-end latency exceeds the human reaction time, the experience is less realistic, with too great a gap between stimulation and response.

Machine-to-machine communications (M2M) is a well-recognized trend. To achieve the system response time required by M2M, the demands on human-to-machine interaction also hold. Industrial robots and vehicles' Electronic Stability Control (ESC) are M2M use cases addressed by the *Tactile Internet*.

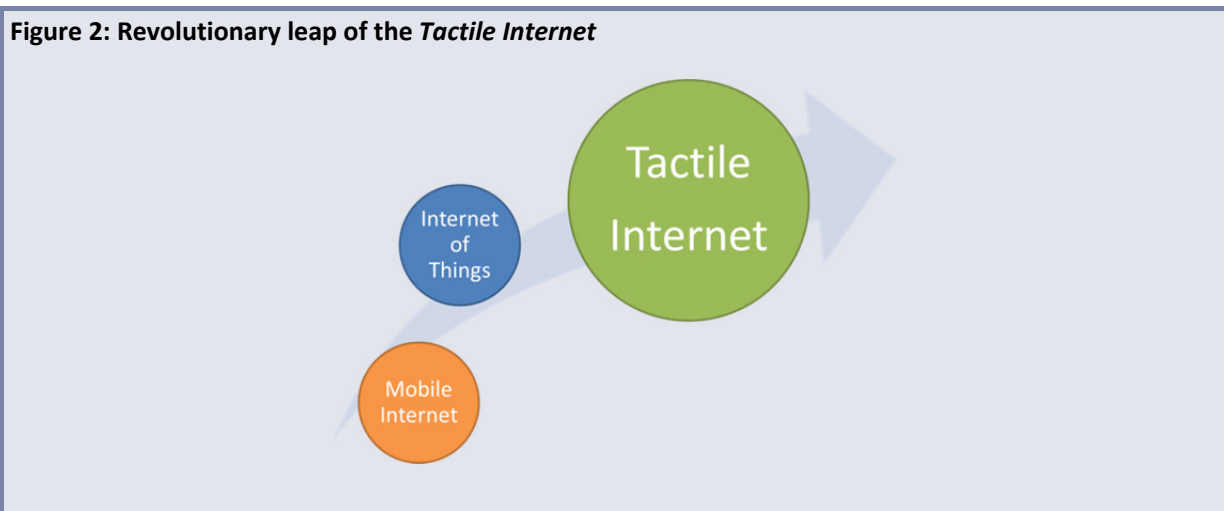
² T. DeFanti und R. Stevens, "Teleimmersion," in *The Grid: Blueprint for a New Computing Infrastructure*, Elsevier Series in Grid Computing, pp. 131-155.

Today’s fixed and mobile Internet infrastructure is typically used for transferring content from A-to-B and is optimized for the transmission of static or streaming content (e.g. e-mail, pictures, voice video). The round-trip latency of communications over existing infrastructure is sufficient for telephony, web browsing and videos with limited resolution.

IoT enables the interconnection of smart devices. Typically, these devices are low-power, resource-constrained sensors with limited functionality designed to transmit low-rate, latency-tolerant data.

The next evolutionary leap is the *Tactile Internet* (Figure 2). The high availability and security, ultra-fast reaction times and carrier-grade reliability of the *Tactile Internet* will add a new dimension to human-to-machine interaction by enabling tactile and haptic sensations. The professional digital infrastructure will similarly revolutionize the interaction of machines.

The *Tactile Internet* will enable an unforeseeable plurality of new applications, products and services. It will enable humans and machines to interact with their environment in real time, while mobile and within a certain spatial communication range. As a driver of innovation, the *Tactile Internet* represents a revolutionary level of development for society, economics and culture.



2. Impact on Society

The instantaneous reaction of the *Tactile Internet* will enhance the way we communicate and lead to more realistic, immersive social interaction in private and business environments. It will also find novel application fields in which to contribute to the solution of the complex challenges faced by our society, challenges including demographic changes (ageing society), the increasing demand for mobility and the transition to renewable energy production (Energy transition or *Energiewende*).

Education and Lifelong Learning: Modern methods for teaching – such as e-Learning, Blended-Learning and Massive Open Online Courses, as well as simulations of movements (e.g. flight simulators) – will be augmented by interactive elements far more advanced than today’s embedded tests or didactic playing elements. The instantaneous reaction times of the *Tactile Internet* will make it possible to enable haptic overlay of the teacher and learner. This will result in novel learning experiences, an example provided by the improvements to be enacted in the training of specific fine-motor skills.

Healthcare: Better quality care will be possible by exploiting the amalgamated expertise of medical doctors connected via the *Tactile Internet* during remote diagnosis and treatment, as well as through the combination of experienced surgeons’ tactile sense with the high spatial precision of robot-assisted operations. Additionally, the support and assistance provided to people with disabilities by exoskeleton-based artificial limbs and power amplifiers will improve their mobility, ensuring them the ability to lead a self-determined life.

Personal Safety Zones: The capabilities of the *Tactile Internet* will permit the creation of a personal spatial safety zone, or ‘bubble’, able to interact with nearby objects also connected to the *Tactile Internet*. Applied to road traffic, in the long term this safety zone will be able to protect drivers, passengers and pedestrians. Vehicles will detect safety-critical situations and react instantly to avoid traffic accidents and warn other objects of impending danger. In production environments, occupational safety levels will improve as production machines or robots will detect and avoid the risk of harm to people in their vicinity.

Traffic in a Smart City: The connectivity of vehicles in the *Tactile Internet* will enable co-operative traffic modes, where traffic flow will be optimized by heeding local safety constraints as well as parameters such as the overall traffic density in a Smart City. Guided autonomous driving or platoon driving will allow for a continuous traffic flow in which safety and energy efficiency will be significantly improved as compared to today’s situation.

Energy: In decentralized electrical energy generation and distribution networks, the *Tactile Internet* enables dynamic activation and deactivation of local power generation and consumption, potentially even taking into account the AC phase information to minimize the generation of unusable reactive power. The *Tactile Internet* is the technical basis for smart grids, providing for improved energy efficiency and stability in electricity networks.

3. Application Fields

3.1 Industry

Automation in industry is a key, steadily growing application field for the *Tactile Internet*. The sensitivity of control circuits when controlling devices moving rapidly (such as industrial robots) requires an end-to-end latency significantly below 1 millisecond per sensor. In typical scenarios of industrial control with closed-loop systems, at intervals of roughly a millisecond, a master station will contact all sensors and actuators and present the acquired data to the control application. In view of the typically large number of sensors (e.g., up to 100 for a printing machine), every individual sensor must be accessed within an end-to-end latency period. There are additional demands for high data rates from other applications; for example in the case of optical sensors employing a high resolution and very high frame rate, such as those needed to monitor the quality of a surface. And when a wireless system is employed, security is of utmost importance as potential attacks could occur wirelessly, without the attacker physically present.

Thus the various control processes – each with specific, deterministic real-time requirements – pose different demands on the end-to-end latency, data rate, reliability and security. In order to progress from specific single-case implementations to flexible, configurable standards-based solutions, both wireless and wired industrial networks should be scalable over a wide parameter range.

Today, control is achieved over a fast wired connection, such as industrial Ethernet. In the future, wired systems are to be augmented or replaced by wireless solutions. The major enhancements required in terms of reliability and attainable end-to-end latency call for huge efforts in research and development.

3.2 Robotics and Telepresence

In recent years, the technical potential of robotics has been demonstrated in various areas including healthcare, autonomous driving, and production and maintenance in particular. Google’s recent acquisition of eight robotics start-ups increased public awareness of robotics’ commercial potential. However, that potential does not come without complexity, and various competitions and challenges have demonstrated the limitations of this technology. For the short- to mid-term future, autonomous robots will only see application in a limited range of rather specific areas, such as autonomous driving.

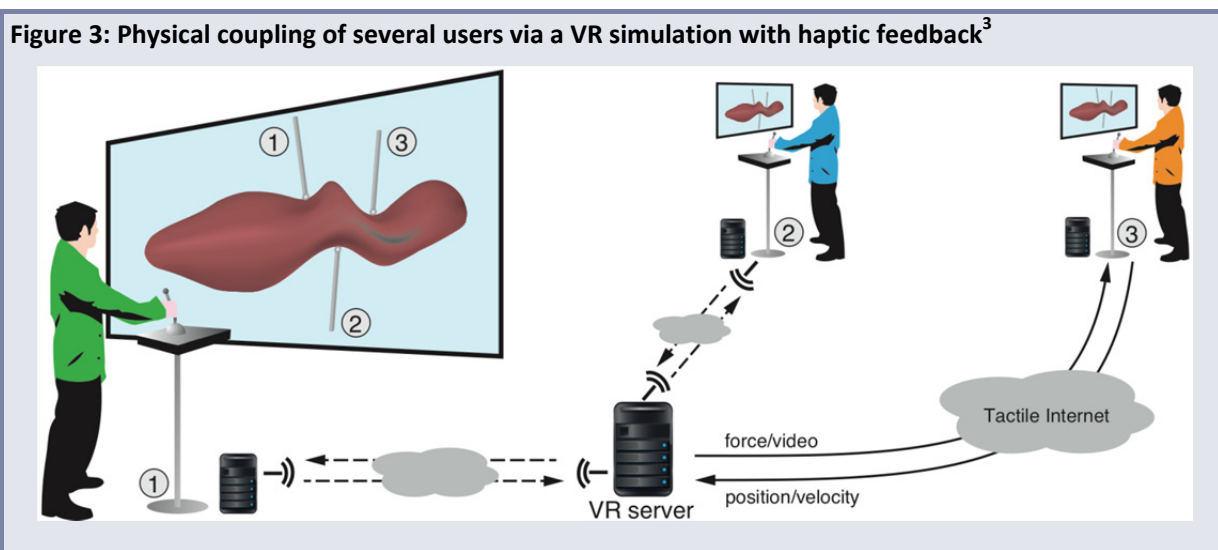
In fields where autonomous robotics is not a realistic option, at least for the foreseeable future, remote-controlled robots are a promising alternative. Construction and maintenance in dangerous areas is an

application area providing a good example of the potential of remote-controlled robots. Such robots have great commercial and societal relevance, as demonstrated by projects such as the planned deconstruction of nuclear plants in Germany, repair work in damaged nuclear or chemical plants, off-shore construction tasks, removal of satellite debris in low earth orbits or maintenance work on the outside of space stations.

A prerequisite for the use of robots in most of these scenarios is a remote-control or telepresence solution with real-time, synchronous visual-haptic feedback. These solutions require system response times (including network delays) of less than a few milliseconds. Communications infrastructure providing this level of real-time capability, mobility support and high reliability and availability is not available today. New technologies and networking concepts will be required.

3.3 Virtual Reality

The *Tactile Internet* will benefit VR by providing the low-latency communication required to enable ‘*Shared Haptic Virtual Environments*’, where several users are physically coupled via a VR simulation to perform tasks that require fine-motor skills (Figure 3).



Haptic feedback is a prerequisite for high-fidelity interaction, allowing the user to perceive the objects in the VR not only audio-visually but also via the sense of touch. This allows for sensitive object manipulations as required in tele-surgery, micro-assembly or related applications demanding high levels of sensitivity and precision. When two users interact with the same object, a direct force coupling brought into existence by the VR and the users can feel one another’s actions. Today’s networked VR systems suffer communication delays too large for stable, seamless coordination of users.

High-fidelity interaction is only possible if the communication latency between the users and the VR is in the order of a few milliseconds. During these few milliseconds, the movements of the users need to be transmitted to the VR server, where the physical simulation is computed and the result is returned to the users in the form of object status updates and haptic feedback. From this perspective, even wired communication may not be sufficient for distant users.

Typical update rates for the physical simulation and the display of haptic information are in the order of 1000 Hertz, which corresponds to an ideal round-trip communication latency of 1 millisecond. A consistent local view of the VR for all users can thus only be achieved if delays are kept very small.

³ Image: C. Schuwerk, E. Steinbach

3.4 Augmented Reality

The development of new augmented reality (AR) applications is accelerating, with the AR market experiencing rapid growth as a result of the availability of small AR glasses and powerful smartphones and tablets equipped with cameras and sensors.

The augmentation of additional information into the user's field of view enables the creation of many assistance systems, for example in maintenance, city or museum guides, driver-assistance systems, medicine, education and the work of police and firefighters.

However, given the limited computational performance of small mobile devices and the presence of considerable delays in network communication, such systems are currently restricted to the display of a small amount of mostly static information. As a consequence, these systems are at this stage only of value to a limited set of applications.

Figure 4: Driver assistance with augmentation of potentially dangerous objects and situations⁴



The *Tactile Internet* will reach new frontiers in assistance systems. Moving from being restricted to mainly pre-processed content, the augmentation of dynamic content or up-to-date information becomes possible. This enables real-time virtual extension of a driver's field of view as shown in Figure 4. Possible dangerous events caused by drivers' not seeing people or obstacles in fog or rain, or as a result of curved roads or car blind-spots, can now be captured by other vehicles able to transmit their view to following vehicles in the form of a projection onto the windshield.

The concept of 'seeing what others see' can also be exploited to great effect in medicine, rescue operations or the virtual attendance of live concerts, as illustrated in Figure 5.

⁴ Image: Modified version of *dangerous driving in the rain - tips* by woodleywonderworks/flickr.com, licensed under CC BY 2.0.

Figure 5: Collaborative capturing of a live event for the virtual attendance of external users⁵

All of these applications rely on the visual combination of real and computer-generated content and require delays of only a few milliseconds in order to avoid irritating the user. This holds true not only for data transmission over fast networks; but also for low-delay sensors and visualization technology; and the new algorithms needed to enable the distributed computation of demanding processes, such as 3D-scene capturing of view correction with low latency.

3.5 Healthcare

Tele-diagnosis, tele-surgery and tele-rehabilitation are just some of the many potential applications of the *Tactile Internet* in healthcare. Today medical expertise is bound to the location of the physician. Tomorrow, using advanced tele-diagnostic tools, it could be available anywhere, anytime; allowing remote physical examination even by palpation (examination by touch). The physician will be able to command the motion of a tele-robot at the patient's location and receive not only audio-visual information but also critical haptic feedback.

The same technical principle of tele-operation makes tele-surgical interventions possible, relieving the patient of costly travel to the surgeon. Future tele-rehabilitation techniques will benefit from the progression of tele-robotic technologies. Envisaged tele-rehabilitation systems see patients wearing an exoskeleton – a robotic device – commanded by the therapist to steer and correct the motions of the patient. The possibility of embedding rehabilitation into the patient's home and everyday life promises higher therapy success rates as well as improved cost-efficiency.

High-fidelity interaction is fundamental to the safe deployment of tele-medical technologies: the physician, surgeon and physiotherapist need a realistic impression of the remote patient's status. High fidelity requires deterministic real-time behavior of the communication channel, which is not achievable with current communications technologies.

Sophisticated control approaches do exist, stabilizing the system in the presence of packet loss and 100-millisecond latency, but the fidelity is significantly decreased. To achieve the high fidelity required for tele-medical applications, it is necessary to achieve end-to-end latency of 1-10 milliseconds and highly reliable data transmission.

3.6 Road Traffic


Mobility is a fundamental need in modern society and crucial to economic development. Road-traffic safety and efficiency are the main factors in sustainable transport. Traffic congestion causes huge

⁵ Image: Modified version of *Garbage live@ KB Hallen #1* by Stig Nygaard / flickr.com, licensed under CC BY 2.0.

economic damage, 17-billion Euros in Germany every year, and the number of vehicles on the road is growing.

In 2011, more than 30,000 people died on European roads. The number of road accidents and fatalities has decreased, at least in highly developed countries, but a considerable and sustainable reduction can only be achieved by vehicle communication and coordination.

Figure 6: Connected vehicles⁶



We have already come to rely on vehicle sensors and driver-assistance systems to support us in arriving safely and comfortably at our destination. Through communication – the data exchange among vehicles (V2V communications) and between vehicles and roadside infrastructure (V2I communications) – a vehicle turns from an autonomous system into a component of a more efficient cooperative one. The data exchange provides information on a vehicle's vicinity as well as non-visible surroundings. Existing communications systems, such as the radio data system in FM radio, bear high latency at low position accuracy and are therefore not suitable for safety applications. Cooperative systems – such as the WLAN-based Car-2-X communications system being introduced to the market – enable a direct data exchange among vehicles but do not support all safety applications.

Applications for vehicle safety require a very low end-to-end latency of below 10 milliseconds, the time needed for collision-avoidance systems to intervene before a collision occurs. With a bi-directional exchange of data for the negotiation of automatic cooperative-driving manoeuvres, a latency of less than 1 millisecond would be needed.

In the future, vehicles will detect a highly dynamic object by radar or video, such as a pedestrian, and disseminate this information to neighbouring vehicles. In the long term, it is expected that fully automated driving will change individual mobility entirely. And with small distances between automated vehicles, in particular in platoons or road trains, potentially safety-critical situations need to be detected earlier than with human drivers.

3.7 Serious Gaming

Serious Games are games designed for a primary purpose other than entertainment. They combine fun with problem-solving challenges and goal-oriented motivation for improvement. Serious Games are relevant to several application fields important to our society; among them education, training, simulation and health.

⁶ Image: Bosch GmbH

Usability and user-experience will be important factors in determining the extent of acceptance and success of Serious Games over the long term. Both demand the selection of appropriate technologies. Novel multimedia and communications technologies will be especially important in the presentation of game content, as well as in the interaction between players and games. The end-to-end delay of communications systems is a key factor limiting the development of game-mechanics, as the perceptible delays directly influence the perceived realism of applications.

Figure 7: Exertion game for adaptive and personalized cardio-training⁷



Exertion games such *ErgoActive* (Figure 7) demonstrate the influence of latency on *Serious Games*: for a high game-quality, a frame rate of 50 Hertz or preferably 100 Hertz is desirable. The control of the games by sensors and actuators should be immediate; actions such as steering, braking or acceleration should be perceived immediately.

The computation of a game’s behavior (game play), including the monitor’s delay and network delays, needs to be less than 20 milliseconds for a frame rate of 50 Hertz and less than 10 milliseconds for a frame rate of 100 Hertz. An end-to-end delay of 1 millisecond is desirable, enough time for the computation of the human’s interaction with the high-quality visualization. This is not a strict threshold; a certain variation of the delay (jitter) is acceptable.

The development of interactive applications has been limited by the high network latency of the Internet. The delays of current technologies are perceptible in games as so-called ‘lags’. This has forced recent games to introduce limits on the number of active players and the distances between them.

Network delays of 30-50 milliseconds or more have been shown to result in a measurable decrease in game-experience and game-quality ratings. Decreasing the end-to-end latency for local as well as long-distance communication will create new, unforeseen opportunities for *Serious Games*. New methods of enabling interaction may include more realistic movements of the whole body and tactile interactions integrated in both single and multi-player games.

3.8 Education and Culture

All of us require fine-motor skills, the ability to elicit precise, coordinated small-muscle movements in body parts such as our hands and fingers. Children having difficulty developing these skills or people needing to regain them in the wake of an injury or illness will seek the help of Occupational Therapists,

⁷ Image: TU Darmstadt

highly qualified teachers or trainers who use hands-on methodologies to intensify the development of fine-motor skills.

The *Tactile Internet* will enable substantially improved learning experiences based on the haptic overlay of the teacher and learner. When the learner executes a task, the trainer will feel the trajectory and degree of applied forces, correcting them as required. Such an overlay operation requires identical multi-modal human-machine interfaces capable of visual, auditory and haptic interaction. Haptic overlay will of course only be possible in the presence of extremely low end-to-end latency between the trainer and the learner. Latency higher than 5-10 milliseconds will prevent the trainer from performing instant interventions to correct the trajectory or degree of the learner's bodily forces.

Challenging, high-risk tasks could be experienced from a safe distance at first, acclimatizing students to the task's demands. Learners could gain a better impression of the necessary fine-motor skills than they would with purely audio-visual information – students at a medical school, for example, could haptically experience a surgery from the eyes, hands and movements of an experienced surgeon at the top of their field.

Similar approaches could revolutionize the process of learning to play a musical instrument. The *Tactile Internet* will allow for remote teaching by an expert; haptic overlay with the student permitting the teacher to perform instant corrective actions. The *Tactile Internet* could also enable music to be played by distributed orchestras. This is impossible today, as high-quality experience will require latencies in the order of a few milliseconds.

3.9 Smart Grid

Efficient, reliable energy transmission and distribution are the foundations of secure energy supply. The increasing usage of renewable energies leads to distributed energy suppliers, which generate energy unsteadily and may inject the generated energy into the power grid at all of its layers. An out-of-phase injection results in 'reactive power', which cannot be used. Energy supply and demand need to be balanced in order to avoid voltage fluctuations. Today's power grid cannot ensure a stable and thus reliable power supply when many decentralized energy suppliers inject power into the grid in an uncontrolled way.

To distribute generated energy, avoid over-capacities and ensure the stability of power supply, smart grids – 'intelligent' power grids – are being developed.

Essentially, a smart grid consists of two components: the power grid, including the generators and consumers; and an accompanying control grid. The smart grid knows the status of power generators, transmission lines and waypoints, as well as the current consumption and tariffs. Based on information on the status of the power grid, intelligent monitors can optimize consumers' power supply and so reduce associated costs. Washing machines and car chargers, for example, will only be activated when favourable pricing is on offer. To stabilize the smart grid, decentralized suppliers will be dynamic, activated and deactivated as required, with synchronous co-phasing of decentralized power suppliers also used to improve stability by power-factor correction.

The technical boundary conditions result in strict latency requirements. Dynamic control that switches suppliers on or off as required will only tolerate an end-to-end latency of some 100 milliseconds. However, a synchronous co-phasing of power suppliers requires an end-to-end latency in the order of 1 millisecond. This 1-millisecond latency results in a phase shift of 18° (50 Hertz AC network) or 21.6° (60 Hertz AC network). Smart grids not only have strict requirements with respect to availability, security and reliability. They also require low latency.

3.10 Additional Application Fields

In the 20th century, the introduction of the assembly line significantly increased the efficiency of the production process and changed the face of manufacturing. This paved the way for the mass production of former luxury goods, putting them in the hands of the general public for the first time.

The challenge today, conversely, is the production of highly customized products. Assembly-line techniques have proved inefficient in this respect, with attempts seeing longer and longer assembly lines developing and expensive robots required only at certain production steps remaining idle for large portions of the production process.

Today's requirements imply a return to the traditional manufacturing scheme. Goods will be manufactured on production islands. Mobile robots will move among these islands, working and delivering assembly parts on demand. The foundation for this is the *Tactile Internet*: a wireless-based, real-time tactile communication network among robots, working on the 1-millisecond latency scale.

Another interesting application domain of the *Tactile Internet* is the use of small un-manned aircraft. Fields of application include the logistics service sector, inventory tasks in forestry and agriculture and surveillance tasks in firefighting, especially in complex environments. Other use cases are the 360-degree recording and transmission of sports events, or providing assistance to police investigations.

Remote control of small unmanned aircraft currently operates in the 2.4 GHz frequency band. This band limits the range of action to just a few hundred meters. The quality of the transmission of video and audio data at this frequency would be limited.

Greater distances, in the order of a mile or more, as well as HD transmission quality, can today be realized with local ad-hoc cellular LTE networks working in the 800-900 MHz frequency band. The end-to-end latency of traditional LTE is low enough to control an aircraft remotely. As with other tactile applications, an end-to-end latency in the order of 10 milliseconds allows control over the aircraft with high precision, without any reaction delay.

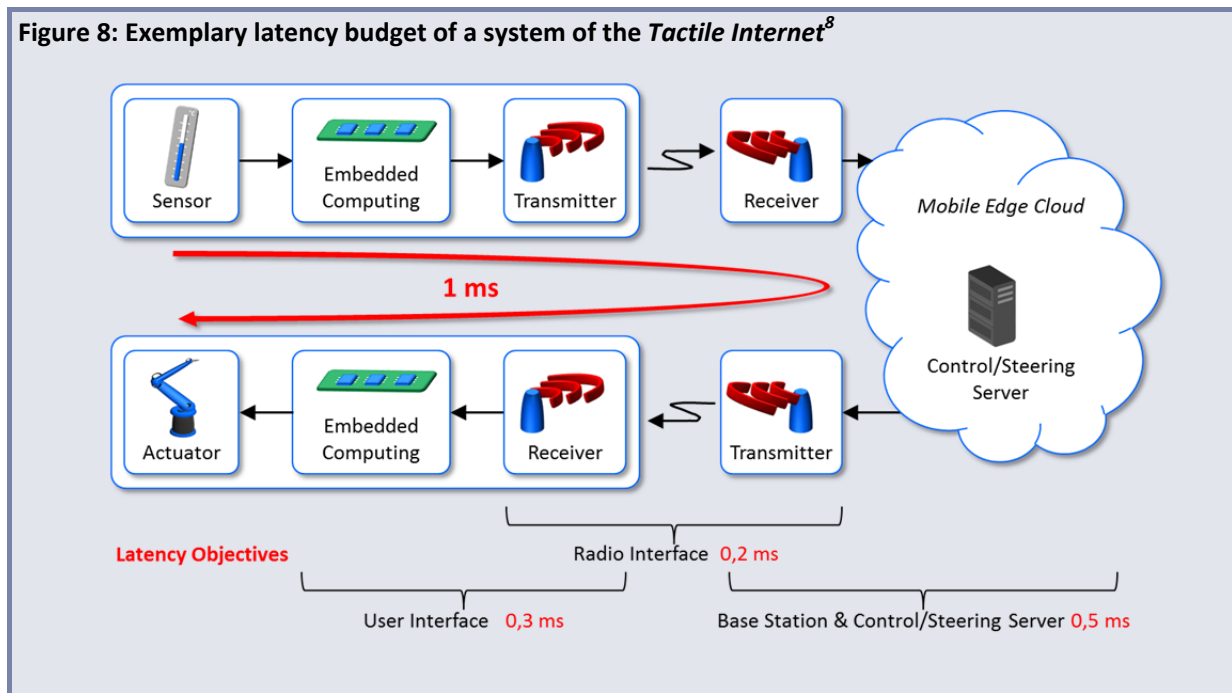
Decreasing latencies into the millisecond region will make it feasible to achieve precise control over a distance of up to a mile. The low latency of the *Tactile Internet* will even enable wireless connections to nearby infrastructures such as data centres, connections which typically require transaction speeds in the sub-millisecond range. The *Tactile Internet* is sure to find further, as yet unknown application domains.

4. Infrastructure Requirements

4.1 Latency and Reliability

In order to obtain 1-millisecond end-to-end latency of the *Tactile Internet*, and thus a system response of 1 millisecond, it is important to understand the chain between sensors and actuators.

Figure 8 shows an exemplary latency budget of a mobile-wireless communication for the *Tactile Internet*. The sensor measures a value. This data is pre-processed and provided to the embedded system controlling the air interface. The air interface then passes the data through all protocol layers to the physical layer. The same happens at the receiving side, for example a base-station with a connected 'Mobile Edge-Cloud' (see Section 3.3.), with the data provided to a control server. It is here that the system control takes place and decisions are made. To obtain the desired system reaction, the decisions are provided to the actuator via this reverse chain.

Figure 8: Exemplary latency budget of a system of the *Tactile Internet*⁸

Beyond low end-to-end latency, high reliability of the entire system is an important quality attribute of the *Tactile Internet*. Demands for the highest possible reliability are associated with requirements for real-time response. This will become clear, with all applications addressed in this section requiring a reliable reception of rapidly transmitted data.

The following sections detail the system components of the *Tactile Internet* and related challenges to be addressed in their development.

4.2 Security

Current communications systems are part of what is commonly known as our critical infrastructure. Our society has high expectations of communications systems, demanding that they be available and dependable, protected from outside attack and free from malfunction. The technical requirements of the *Tactile Internet*, especially ultra-low end-to-end latency, pose enormous challenges for communications systems. Paradigm changes are necessary, to ensure both data security and the availability and dependability of systems.

In today's communications systems, a secure communication is based on separating the encryption from the transmission technology. Classical encryption methods provide security against eavesdropping only when the encryption algorithms are sufficiently complex and the eavesdroppers are limited in their computing power. Because these security mechanisms can only be implemented at higher protocol layers, this leads to noticeable delays. The task of the transmission technology is purely to provide high transmission rates.

For the *Tactile Internet* to provide securely transmitted data with very low end-to-end latency, the security of communications against eavesdroppers and attackers must be embedded in the physical transmission. Suitable coding techniques will ensure that only legitimate receivers are able to process a secure message. The nature of a secure message is such that an eavesdropper, as an illegitimate receiver,

⁸ Image: I. Riedel, G. Fettweis

cannot decode the data. Even with infinite computational power, an eavesdropper will not be able to decode the message. In a mathematical sense, this approach provides absolute security.

The important criteria in assessing the performance of new communications systems are, among others, their maximal rates of transmission of absolutely secure messages as well as their maximal key-generation rates of absolutely secure keys and rates of transmission of absolutely secure keys.

This approach – prioritizing the technical security of transmission – is superior to classical encoding methods in that physical propagation conditions directly determine the performance variables of communications systems. This permits system dimensioning and optimization to be taken into consideration.

Another central challenge faced by the *Tactile Internet* will be the identification of users, or terminal devices to be specific. The existing method of separating authentication and physical transmission does not allow for low end-to-end latency. The *Tactile Internet* thus necessitates that authentication be embedded in the physical transmission. Hardware-specific attributes of the transmitter can be used to achieve this, such as forms of biometric fingerprints known as ‘Physical Unclonable Functions’.

4.3 System Architecture

The *Tactile Internet* requires the best possible response times, availability, reliability and security. These objectives can only be accomplished by a distributed service platform architecture. The need for ultra-low end-to-end latency requires that tactile applications be kept local, close to the users.

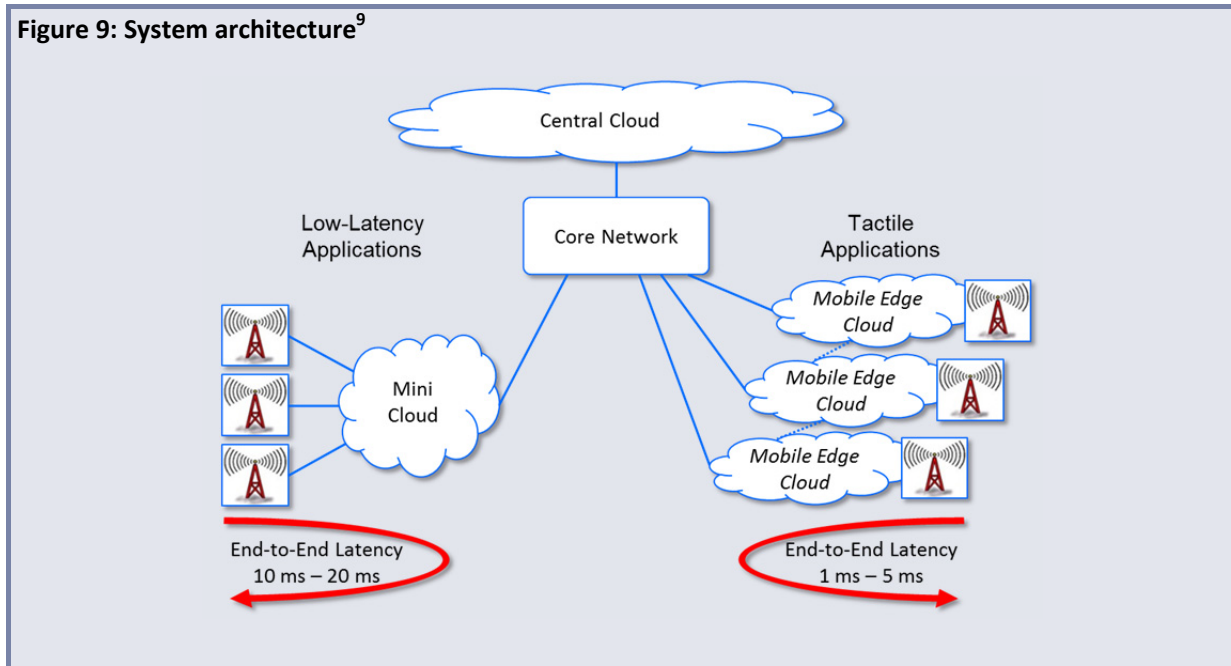
In an initial step, a small local data center or ‘Mini Cloud’ may support a cluster of user devices in its vicinity. This Mini Cloud then provides the complete functionality of a wireless network as well as all tactile services required within the cluster. Such a Mini Cloud could control and steer all of the robotics of a large industrial complex.

The next step would be to deploy a Mobile Edge-Cloud responsible for the tactile applications’ network functions at the edge of the mobile access network, close to the user devices.

Mini Clouds and Mobile Edge-Clouds include three basic functional units:

- (1) A cloud-based service platform for tactile applications, online games, web caches and other delay-sensitive services.
- (2) Virtualized network control functions, performance optimizers and security systems.
- (3) All necessary interfaces with the wireless network.

Step-by-step, a multi-stage hierarchy of cloud platforms will arise: many Mobile Edge-Clouds at the level of the user devices, Mini Clouds at the local level, and a limited set of larger central clouds.

Figure 9: System architecture⁹

4.4 Sensors and Actuators

Communications involving sensors and actuators is associated with extremely demanding requirements on both hardware and software. To build a reliable control system, well-determined response times must be guaranteed. Some applications permit end-to-end latencies of no more than a few microseconds: the time for every sub-step (data processing in the sensor, communication with the control processor, generation of the control signals, and transmission to the actuator) must lie below this upper limit.

The first step in the processing chain is signal-conditioning in combination with data-compression. Intelligent compression methods, which significantly reduce the amount of data to be stored or transmitted, play a central role in system optimization. The compression must be done with minimal algorithmic delay and should essentially work without buffering. In recently developed event-triggered control systems, sensor data is only transferred if new significant information is present. Plausibility estimators are also used to eliminate outlying values. The combination of modern filtering techniques with a suitable physical model allows for a good predictability on the basis of only a few measured values.

Promising solutions are based on *'in-network processing'* methods. These methods enhance the raw data delivered by the sensor node with additional information at every node between the source and the destination. In this way, a value processed by the control processor has a much higher significance, leading to more efficient, higher quality control decisions.

Future networks will require major improvements in the context of inter-node communications, since the delay caused by medium access often lies in the range of a few milliseconds. In such a situation, sensor data could be stale even before it is transmitted. To reduce the delays at all protocol layers, innovative and scalable procedures for protocol design are needed. A decisive role is played by the end-to-end latency from sensor to actuator, which includes all intermediate steps. An intelligent coordination of the relevant components is a good way in which to reduce the latency. Here it should be recognized that the individual control processes are strongly coupled, mandating a high-level analysis of the system over and above the level of separate processes.

⁹ Image: I. Riedel, W. Häffner

In addition to fast compression algorithms and short transfer times, minimizing the end-to-end latency requires highly efficient signal processing at the user interface. Specifically, in the case of visual interfaces, the *Tactile Internet* needs immediate image processing, increased frame rates, and high communication data rates. The final goal is to bring the latency of digital video systems down to that of analog solutions, for which the limiting factor is the frame rate. This aspect is of high importance for applications which must satisfy strict safety standards, such as driver-assistance systems.

With haptic user interfaces, the upper bound for tolerable delays is even more stringent. Immediately after a sensor value has been obtained, it must be transferred. Only under this condition can hard surfaces be sensed haptically, in a realistic way. The stringent delay requirements call for new approaches to the processing and compression of haptic data streams. A system response time above 10 milliseconds would lead to instabilities in haptic interaction.

4.5 Access Networks

The *Tactile Internet* will set demanding requirements for future access networks. The outlined use cases will require round-trip latencies of as little as 1 millisecond as well as high reliability and capacity (data rates). In some cases, wired access networks are partly meeting the requirements, but wireless access networks are not yet designed to match these needs. Scaling-up research in this area will be essential.

Future access networks need the ability to cope with one-way latencies of only 200 μ s (Figure 8). The resource allocation of the available physical blocks needs to be done up to 10-times faster than in LTE. VR applications will require haptic feedback that demands large volumes of data being transmitted reliably and ultra-fast.

The transmission errors inherent to wireless systems will need to be ironed-out through careful design. Applications in industrial environments for example, where potentially huge numbers of robots and machines will work in close proximity, will create challenging interference conditions not satisfied by current wireless systems. Classical approaches for medium-access control are not suitable in such an environment, and new techniques are needed to drive latency down to a bare minimum.

New ideas and concepts to boost access networks' inherent redundancy and diversity need to be researched to address the stringent reliability requirements of *Tactile Internet* applications. Simplified and fast resource-access schemes as well as efficient signaling protocols need to be designed to optimize the use of the underlying physical resources.

4.6 Mobile Edge-Cloud

Light travels 300 kilometers within 1 millisecond. The distance between a control server and the point of tactile interaction can thus be 150 kilometers, at most. This distance assumes no processing delays in the communication path. When the signal-processing, protocol-handling and switching delays are taken into account, it is clear that the control server needs to be relatively close to the point of tactile interaction. A Mobile Edge-Cloud could process the more latency-critical server requests, with less latency-critical server requests handled by a server located further away in the cloud (Figure 9).

The *Tactile Internet* comprises multiple components. In addition to the control server of the Mobile Edge-Cloud, the control system will include the embedded systems of the mobile sensors and actuators. For this reason, a distributed control system is generated. The software running on the control server must of course meet the latency requirements, with the control system demanding a real-time operating system that directs incoming control information as fast as possible to the control software to guarantee rapid response.

In order to achieve the high reliability of the *Tactile Internet* without increasing latency, multiple transmissions over parallel communication channels are required. Future operating systems of the communications infrastructure will monitor and manage the service quality of the communication.

A seamless user experience in today's wireless communications systems requires a 'hand-off' – one base-station hands the communication link of a mobile user over to another base-station. In addition to this communications hand-off, the hardware and software system running on a control server of the *Tactile Internet* must also be built to hand-off the context of a running application from a server at one base-station to a server at another.

The particulars of this hand-off and the *Tactile Internet's* requirements for robust parallel communication and real-time operating systems are completely new challenges, without existing technical solutions.

5. ITU Framework

Figures released by ITU in May 2014¹⁰ indicate that, by year-end 2014, there will be almost 3-billion Internet users, representing 40 per cent of the global population; two-thirds of them in developing countries. The number of mobile-broadband subscriptions will reach 2.3 billion, with 55 per cent of these subscriptions in developing countries. Mobile-cellular subscriptions will reach almost 7 billion; the increase mostly the result of growth in developing countries, where mobile-cellular subscriptions will account for 78 per cent of the world's total.

As impressive as these figures seem, the number of connected people will soon be dwarfed by the number of connected *things*. Two oft-quoted statistics in relation to the forecast growth of IoT are Cisco's prediction that there will be 10-billion Internet-connected devices by 2018¹¹, a figure which includes M2M modules, and Ericsson's vision of more than 50-billion connected devices by 2020¹².

For all and everything to be connected, international standards to define the common communication protocols and functional frameworks will be fundamental, demanded by all industry segments. Without standards, fragmented markets will experience a lack of interoperability and scalability and costs too high to achieve significant growth.

Innovation in the organization of international standardization work will be crucial in meeting the new challenges to arise in the future. The ICT standardization landscape has become a kaleidoscope of various stakeholders, with ICTs now crucial to every industry sector and the delivery of essential services to the public. Standardization processes are expected to evolve in line with the diversity of the Tactile Internet's application areas and associated stakeholder ecosystem. The collaboration and cooperation of these stakeholders is indispensable.

Wireless communications rely on a common resource: Frequency spectrum. Good spectrum management is a prerequisite for the development of wireless communications and, much like ICT standardization, it must preempt technological advance and thereby offer its supporting framework.

Spectrum management is based on three pillars, all equally essential: spectrum allocations, spectrum management and spectrum monitoring. Sound spectrum allocations require extensive international cooperation to ensure that regional, or preferably worldwide spectrum harmonization takes place to deliver the benefits of economies of scale and international roaming in a way that encourages investment in the development of new technologies without threatening past investments in legacy wireless communications networks.

¹⁰ Source: ITU: The World in 2014: ICT Facts and Figures. May 2014, <http://www.itu.int/en/ITU-D/Statistics/Documents/facts/ICTFactsFigures2014-e.pdf>

¹¹ http://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/white_paper_c11-520862.html

¹² http://www.ericsson.com/ericsson/press/events/2009/business_innovation_forum/Hakan_Eriksson_Ericsson.pdf

Given ITU's vital role in the global management of the wireless frequency spectrum and satellite orbits, its global reach including developing countries, and the wealth of experience in defining new-technology landscapes through the publication of the requisite standards, ITU's unique public-private membership model offers an ideal platform for collaboration open to all invested in realizing the vision of the *Tactile Internet*.

ITU-T Technology Watch surveys the ICT landscape to capture new topics for standardization activities. Technology Watch Reports assess new technologies with regard to existing standards inside and outside ITU-T and their likely impact on future standardization.

Previous reports in the series include:

Big Data
Spatial Standards
Mobile Money
Smart Cities
E-learning
E-health Standards and Interoperability
Privacy in Cloud Computing
Digital Signage
Trends in Video Games and Gaming
The Optical World
Decreasing Driver Distraction
Biometrics and Standards
The Future Internet
Distributed Computing: Utilities, Grids & Clouds
NGNs and Energy Efficiency
Remote Collaboration Tools
Ubiquitous Sensor Networks
ICTs and Climate Change

<http://www.itu.int/ITU-T/techwatch>