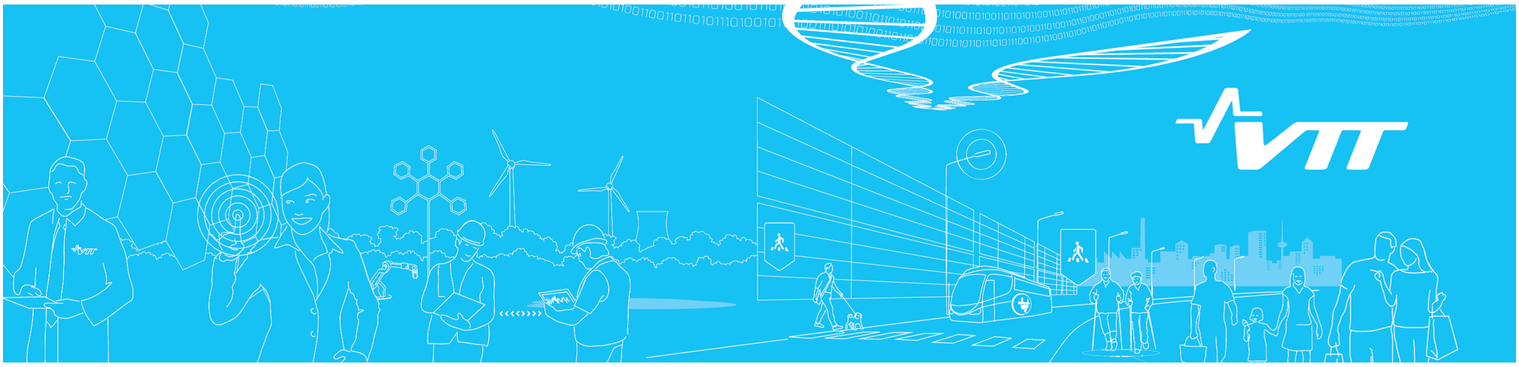

WHITE PAPER

Human Verifiable Computing in Augmented and Virtual Realities

Authors: Kimmo Halunen, Outi-Marja Latvala, Hannu Karvonen,
Juha Häikiö, Seppo Valli, Maija Federley & Johannes Peltola

Preface

> *Trust is the glue of life. It's the most essential ingredient in effective communication. It's the foundational principle that holds all relationships.*
>
> — Stephen Covey

In this white paper we present some of the results of VTT's iBET project *Human Verifiable Computing* (HVC). The research problem tackled in the project was, how to provide users with reliable information about the trustworthiness and security of the digital systems they are using. In our work, we decided to utilise augmented and virtual reality (AR and VR) technologies to display this information to the user and to see what types of challenges these technologies pose for HVC.

Solving the problem of HVC would have many benefits not only in AR & VR, but in any system that outsources computations from the user to a digital system. HVC could bring us ways to establish trust to artificial intelligence systems and towards robots and other autonomous systems, e.g., cars and ships. In this sense, HVC is an important part of our digital future where human-computer interaction is becoming increasingly seamless and essential part of our everyday activities both at work and leisure time.

Although we have made some advances in solving the problem, there are still many open questions and challenges that need solutions. The main purpose of this document is to introduce these research questions to a wider community of researchers and practitioners. This way we hope to find common topics of interest, where we can work together towards the goal of Human Verifiable Computing.

The problems span over several disciplines and there are theoretical problems, engineering challenges and also issues related to psychology and human behavior that need to be solved. Thus, we hope that many people and organisations working on these disciplines will join our effort.

20th November 2017

Authors

Contents

1. Introduction

Augmented and virtual reality (AR & VR, respectively) are quickly becoming commonplace in many areas of life. Although first considered in domains such as gaming and other entertainment, new applications emerge in various areas of life, such as training and maintenance work in industries. AR and VR technologies also bring new directions to cyber security of systems, and their security and reliability become increasingly important when applied in more safety critical domains.

Machines and digital systems are increasingly in interaction with human users and there is a trend towards more natural forms of interaction such as speech and gestures. In addition, the ways in which these digital systems communicate with users are changing and many devices do not necessarily have displays and keyboards like our traditional computers. Thus the old ways of providing information to the user, such as pop up messages in textual form, are no longer useful or even feasible at all.

In our opinion there is a need to design and build systems that can communicate the trustworthiness and security to the user in a way that the user can intuitively understand. We call this concept Human Verifiable Computing (HVC) and during our project here at VTT, we have produced some initial steps in solving the problem through AR and VR technologies.

In this paper, we describe the problem of HVC and how it could be approached. We give some initial ideas for potential solutions, but the main contribution is in pointing towards research challenges that need to be solved in order to truly realise our vision. We invite all interested parties to continue research on this topic.

1.1 What's the Problem?

The problem statement that we have set out to solve in a VTT project named "Human Verifiable Computing" is fairly simple:

> *How can we present the user with reliable information about the trustworthiness and security of the AR/VR system in a human understandable form? And how can we prevent this information from being corrupted by an attacker?*

As with so many other fairly simple questions, this question does not have a simple answer or a solution. Human senses are easily fooled and in a AR/VR setting, the system that generates the inputs to our human senses could be under the control of a malicious party. Is there anything that could be done to verify the correctness of the system? There are similar challenges, i.e. a need for transparency and verifiability, also in other domains such as artificial intelligence and cloud computing.

As it turns out, there are many ways to measure the trustworthiness and correctness of a software or hardware system. The main shortcoming of these is the fact that the results of these measurements are not comprehensibly communicated to human users and the user is left to trust the designers, manufacturers and vendors of their systems. Thus, there is a gap between the machines and the human user. This gap is what we want to address with our research.

2. How to Approach This Problem?

How can we approach this problem? Is it not enough to trust the designers, manufacturers, service providers and other parties that provide the needed functionality? Unfortunately, there are so many different interests at play, that for any given user it is not possible to simply trust that all involved parties act with the user's best interest in mind.

There are many ways in which humans interact with computer systems and this research can bring some clues on how to proceed. It is important to note, that there are many ways in which human users are authenticated towards various computer systems (see for example [5] for a survey on user authentication mechanisms). The other way around is still missing. That is: "How can users authenticate or verify the dealings of the computer systems that provide us with services and information?"

In general, the advances in artificial intelligence (AI) and machine learning are mostly aimed towards helping the machines to understand and interpret humans. This can be seen from the many inferences that our social media interactions reveal from us [9] and how for example machine learning can learn emotions and other data from micro expressions [21].

The ethics discussion around this topic is also interesting, but it has not given tools for building the kind of trust we are looking for [2]. The focal points there are how to have AI behave ethically, the possible privacy implications of large scale data analytics [3] and how to prevent AI (algorithms) from learning and amplifying biases currently manifesting in our societies. Furthermore, there is interest in having the manifestations of AI, such as robots more closely comparable with humans. A recent development towards this is the granting of citizenship to a robot in Saudi Arabia.

Despite these developments, we are not building methods that enable human users to easily understand how different computations and systems such as AI work. Even worse, there are very few methods that even try to solve this and they are very limited in their scope [4, 14]. In our view, this needs to change and we envision that augmented and virtual realities could provide tools in tackling this problem. However, AR and VR are not the only possible technologies to tackle this.

2.1 Augmented Reality

Augmented Reality (AR) is used for superimposing virtual objects in the user's view of the real environment. The real environment (real world), and a totally virtual representation (VR) are the two ends of the Mixed Reality (MR) continuum, augmented reality being situated in the middle of this continuum [15]. AR technologies may be used to provide novel visualization functionalities for a wide range of applications.

In AR, an essential challenge is an accurate and fast enough mapping or linking between a physical environment and a visual representation (e.g. still image, sequence of images, video, 3D model, or 3D animation). Mapping is generally assisted by adding a priori information to the view in the form of visual markers, which, however, may be disturbing to the eye.

Output risks are present in any digital service visualizing data to a user, especially when it comes from several sources or even from competing applications. These risks are related for example to overloading of the user by information, being partly non-relevant or even spam, or relevant information being occluded or hidden behind bigger or less relevant overlays. In

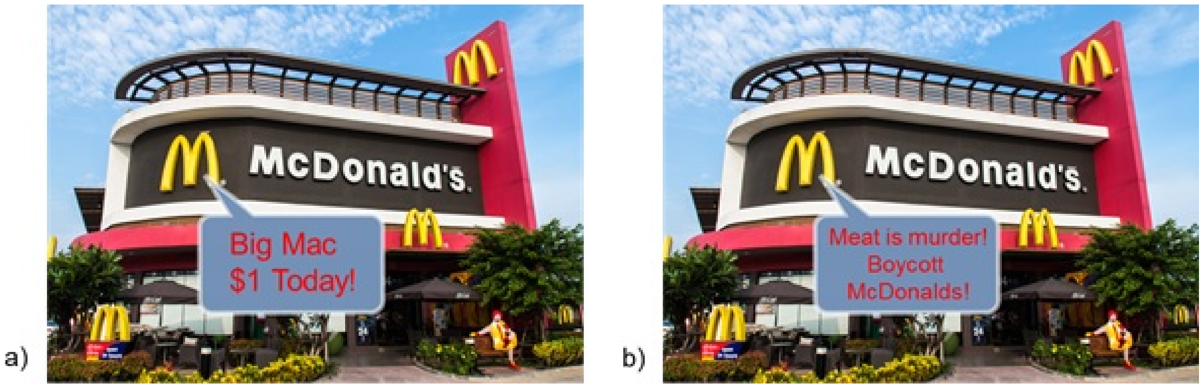

Figure 1. Illustration of a malicious augmentation in world browsing AR service: a) what we expect; b) what we might get.

the paper by Lebeck et.al [10], output security management was listed to give answers to the following type of questions (direct quotation of [10] is shown in italics):

- *Who (i.e., which application) displayed particular content. Knowing this could, for example, be useful in disambiguating content generated by a phishing application or advertisement from content generated by a legitimate banking or route guidance application.*

- *What kind of content a particular application can draw. For example, should an automotive application be able to draw virtual pedestrians on the road?*

- *When an application can draw, based on the context of the user's actions or environment. For example, could an HMD texting application pop up a full-screen message when the user is doing something potentially hazardous, like walking down stairs?*

- *Where an application can draw, both on the display (i.e., with respect to the user's screen") and within the world (i.e., with respect to specific objects or 3D regions in the world). For example, could an automotive application render an ad on top of a road sign?*

The relevance of output risks is bigger when the user's awareness is not properly supported to make a distinction between different information sources. This occurs especially in AR services, which overlay virtual information on real-world views, thus making it more difficult to distinguish between the two realms. Paradoxically, still at the moment, the lack of photorealism or fidelity of virtual elements helps to make the distinction. By more seamless augmentations, the problem gets even more relevant. This is one issue, which points towards the importance of having human verifiable visualisations in AR.

Many times the risks are related with both input and output, so that the classification to either of them is somewhat artificial. For example, when a particular set of captured features (cf. input) is used for rendering a malicious or hostile augmentation (cf. output) on a company facade. Figure 1 illustrates a case possible in an AR service for world browsing.

The example in Figure 1 refers to big unsolved questions relating AR services: Who owns the visual appearances of properties or items? Does any law protect against malicious virtual augmentations on recognized objects, or even identified faces? How to protect against this kind of frauds? Copyright, property, privacy etc. laws do not likely protect visual appearances as they do physical properties and individuals. The above examples show that there are even needs for new laws relating privacy and security in AR services. An interesting question is also, to what extent technical solutions can prevent such unauthorized augmentations. Note that some discussion on legal issues is included in [22].

2.2 Virtual Reality

Virtual reality (VR) views are typically complete renderings by one service provider (for example a virtual world service or a building model). Thus, communication about security for the user can be incorporated in the design of the environment and managed by the service provider without needing to take into account potential input risks. In VR, privacy is typically not violated by capturing pictorial data from the users' environment. Privacy is more about knowing higher-level context of the user – i.e. that he/she in general is using the service. Respectively, security is more related to access rights to the service content. Authentication and access rights are rather well solved in VR services. This is partly due to that VR technologies and services are generally more mature than those for AR, but also due to that VR does not need so much data from users' environments as AR does. However, capturing user's motions by a Kinect type of sensor for animating a VR character or an avatar poses a privacy threat – even if not necessarily revealing user's real environment.

A virtual reality scene may provide different views and interactions to different users with varying access or authorization levels. Different access levels may show up as different renderings of the complete 3D modelled information for each user. This may be the case for example when visualizing the properties of a 3D modelled building for a diverse group of stakeholders or visitors. Thus, information about user rights and being able to manage them could be useful for an administrator.

Furthermore, it is again of paramount importance that in VR (as in AR) the user can trust that the virtual reality corresponds to what the user experiences in the VR. There are for example fully fledged payment terminals embedded in VR that could be applied in VR with the user's own credit cards. And there are of course many other scenarios that require trust between different users and the VR engine and service providers.

2.3 Verifiable Computing

In cryptography, verifiable computing means that a client requesting an external party to carry out a computation for it gets a verification that the computation was carried out correctly. The verification can happen *online* and require interaction between the parties. The other possibility is to provide the client with a verification "tag", which the client can verify offline, without interaction to the external parties. These verifications should have much lower computational cost than the original outsourced computation in order to be useful. An excellent summary of verifiable computing can be found in [27].

However, these technologies only provide assurance between machines, like many other cryptographic schemes. This means that the human user is again left to trust the correct functioning of their devices and software. The results of verifiable computing or the protocols for checking the computations are not accessible to humans.

There are some systems that aim towards having the user included in the verification of security and trust. PRISM [4] is a system for establishing trust to a simple device through human means. However, the limits that the PRISM requires of the system are very strict and in essence unrealistic for a modern computing device. The method itself is fairly simple for a human user, who only needs to check a list of challenges, provide the response and measure the time it takes for the device to respond.

iTurtle [14] is a method for establishing trust to a device and/or the software that it is running. The authors envision that this type of device could be used as a root of trust to further interactions with digital systems. The authors do not specify a concrete example, but they do state that the status of the device could be displayed to the user via a led or some other simple visual cue.

3. Human Senses and Computer Feedback

As desktop computers and smart phones with a screen are so ubiquitous nowadays, the traditional way computer systems provide information and feedback is visually through a graphical user interface. Typically, also auditory feedback can be given to capture the user's attention.

However, recent advancements in computing have enabled systems to be embedded increasingly into our everyday objects, such as cars, home appliances, clothes and other accessories. For example, a modern car can contain millions of lines of code in its software system. However, when driving a car, the visual modality is already reserved most of the time for the primary task of driving and the possible user operation of in-car computer systems is usually a secondary task.

Similar issues can be found when considering humans moving in their daily environment where different type of relevant data can be embedded. For example, augmented reality glasses may provide to some extent visual information for the human, but it may be disturbing the visual field in some critical tasks.

Therefore, the researchers and developers of modern computer systems have started looking at the other human sensory modalities to as primary sources for natural feedback for the user. For example, we can see the rise of user interfaces where auditory (hearing), tactile (haptical), or even olfactory (smell) sensory channels are utilized as the primary way of feedback.

Next, we discuss the human sensory modalities in more detail and discuss how they have been applied to feedback given by computer systems.

3.1 Vision

In modern systems our visual modality of interaction is utilised the most. Text, pictures, symbols and video are presented to the users in almost all digital services that they use. In addition, both AR and VR rely heavily on visual augmentations or interaction with the users.

Vision is thus important and it should be used to great effect also in HVC. In our work, we have used visual cryptography [18], where the user can decode a cryptographic message by just looking at the (correctly formend and positioned) encrypted messages. There are also visual hashes, where user is asked to check if two images are the same, before confirming some action. This is similar to comparing numerical values or other codes for similarity.

However, as we all know, human vision and visual recognition is far from perfect. People make mistakes constantly, but it is usually easy to recover because we also continuously re-evaluate the information presented to us by our vision (and compare it with other senses). This is something that our current digital systems are not very good at facilitating. Many times trust and security related decisions and visual cues and comparisons are made only once or the re-evaluation is done only after considerable time. Thus, we think that the way vision is used in HVC needs to be geared towards more continuous methods in order to be more useful and intuitive to users.

3.2 Hearing

Human hearing (or auditory modality) is also used in our interactions with digital systems and also in AR and VR. Hearing is not utilised as much as vision and it is definitely very rarely used to communicate security or trust information. In other contexts, such as fire alarms and emergency vehicles, sound is used both more often and to a greater effect.

There are many qualities in human hearing that suggest that it could be more extensively used also in HVC. Humans are usually quite keen on picking up differences in audio and it is actually quite hard to synthesize a convincing audio fraud, if the sound is familiar enough. This could make audio a good venue for transmitting also security related information.

In addition to the actual sound or voice, one can also use the content of the sound (for example speech) as a secondary factor. The content could also utilise such personal data and cues that might be hard to guess or know by attackers. For example, informing of a meeting location by "Let's meet at the same place as last week.", requires the attacker to gain the information about the place and not just access to the audio.

On the other hand, using audio requires capturing the attention of the user and making sure that the message is not drowned in the possible noise that surrounds the user. When the user is using headphones or other such device, this can be alleviated.

Sound can also be used to convey information about the direction of an event (if such information is meaningful) and also on the magnitude (loud vs. soft sounds). In addition, the type of the voice can be used to inform the user. It is also important to note, that sound can be used for interaction for example through speech recognition.

3.3 Haptics

An origin of word haptics is related to the Greek word haptestahai, which means relating to or based on the sense of touch. Thus haptics means our tactile modality of interaction. Touch is an essential sense for humans in daily life and it has been utilized in different types of haptic user interfaces in different domain areas, such as entertainment, assistive technology, automotive manufacturing/assembly [6]. Haptics can be studied from multiple different perspectives and it can be divided into 1) Human haptics that refers to human sensing related to touch, 2) Computer haptics that refers to software regarding touch and feel of virtual objects, and 3) Machine haptics that refers to design and use of machine, which can augment or replace human touch [25].

Sense of touch is typically utilized in parallel with other sensory channels, typically sighting and hearing, for creating solutions for multimodal interaction and taking fundamental human capabilities into account [23]. Haptic feedback channels can be divided to the following categories; 1) tactile sensations including e.g. pressure, texture, puncture, thermal properties, softness), 2) vibrotactile sensations related to the oscillating objects in contact with the skin, and 3) kinesthetic perception relating to awareness of a body state [7]. Haptic technologies can be classified to wearable cutaneous devices, active surfaces and mid-air haptics [19]. They all have their special characteristics and targets of use.

During recent years, there has been a growing interest in research and industry to use of haptics in virtual (VR) and augmented reality (AR) based solutions. Research on haptic feedback is critical for achieving more natural and intuitive user experience in VR environments through making it possible to touch virtual content directly with hands [19]. Currently, main limitations of haptic systems are related to capability to render a desired force or impedance in different contexts [25]. Recently, research on haptic feedback and virtual spaces in combination with fast connectivity technologies has raised interesting emerging opportunities [1, 8]. In addition to haptics related technology oriented studies, research should focus on user-centered research including user's cognition and emotional experiences in the future [24].

3.4 Other Senses and Multimodal Feedback

Although the three sense mentioned above are definitely the most mainstream, we should not restrict ourselves to only considering those in our interactions. There is a very interesting research area on utilising the sense of smell (olfactory modality) in VR by synthesising smells from a small number of basic odors. There is for example an ongoing project at University of Tampere that aims for real-time syntehsising of scents [1].

Also our sense of balance (vestibular modality) could be used to inform users about trust and security of a system. This type of information is not something that could be used in all scenarios, but in VR this type of feedback to the user may be desirable in order to create better immersion. Thus, it could also be used for informing the user about potential problems in the system.

One of the most important issues is, that there is a significant number of people that have at least one of their senses somewhat impaired or even lacking completely. Thus, relying on just single modality to convey such critical information as trust and security is not useful. HVC needs to be accessible to all people and thus multimodal feedback is a necessary feature of such system.

This multimodality should also be fluid and context-sensitive. That means that the modalities used to represent HVC information can change depending on the environment and possible task of the user and also on the capabilities of the user herself. In a loud environment sounds might not be used or only used to convey very simple information (like smoke alarm) and if the user is engaged in a task that requires her visual and auditory faculties, maybe only haptics or smell should be used.

The possibilities of multimodal feedback require more research not only in the context of HVC but also in general. We think that such technologies will be useful in HVC and that they should be utilised where possible.

4. Possible Solutions

What could be the possible solutions to the problems that we face with HVC? First, we need to understand how trust is formed between human users and different machines and digital services. Then, we should also see how modern cryptography can be utilised in building this trust and communicating the correct information to the users. Here we present some of our initial solutions and directions of research.

4.1 Building Appropriate Human Trust in Automation

From a human factors point of view, trust is a psychological concept. Several studies (e.g., [11, 13, 16, 17, 28] have shown that an operator's (i.e., user's) trust in automation (i.e., com-

[1] see `http://www.uta.fi/sis/tauchi/esc/projects.html` for more details

puters) is one of the most important factors affecting the usage of an automated system. Automation can be here defined as computing technology used to automate tasks that have been previously conducted manually by a human operator. Automation is nowadays used almost in every technological environment from cars to nuclear power plants. In this context, trust can be defined as 'the attitude that an agent will help achieve an individual's goals in a situation characterized by uncertainty and vulnerability' [12]. The 'agent' in this case is automation, which works in active collaboration with other parts of the system on behalf of the operator.

In general, if trust in some particular situation is needed, the nature of this situation must be uncertain and something valuable must be at stake. This valuable can be related for example to safety (e.g., a dangerous threat), productivity (e.g., economic losses), or health (e.g., an injury) [26]. Therefore, trust in automation is an especially significant factor with safety-critical systems as there can be fatal consequences if something goes wrong. Trusting the automation means for example that the operator sees that the risk related to using the automation is lower than conducting the same task manually.

According to [12], automation's capability affects considerably the development of operator's trust in automation. However, [12] does not strictly define what they mean by automation's capability, except 'trustworthiness'. We see that automation's capability refers both to the scope of technical possibilities of the automa-tion system and to the technical reliability of the automation. Therefore, automation's capability basically means whether the automation is capable of finishing the given task correctly. Other factors affecting the development of trust in automation – but not discussed here in more detail – are for example the provided user interfaces (e.g., how transparently reliability-related information is visualized), the amount of correct and false alarms by the automation, operator's professional self-confidence, fatigue, workload, task complexity, the provided training, and cultural/organizational background.

The amount of trust can define whether the operator is using the automation somehow inappropriately. Lee and See [12] describe two possible reasons for inappropriate use of automation: overtrust and distrust. Overtrust means that the operator perceives the automation's capability to be greater than it actually is. When this occurs, the operator continues using the automation even though it does not work the way it is supposed to work in that situation. As a consequence of trusting the system too much, the operator might not monitor the automated system on a sufficient level or cannot recognize the system's restrictions in a way that should be needed.

Distrust on the other hand refers to a low level of trust compared to the actual capability of the automation. Because of distrust, the operator might not utilize some parts of the automation or might not use it at all. When an operator's trust in automation is at the same level as the capability of the automation, it can be said that the trust is at an appropriate level (well-calibrated). Therefore, the operator uses the automation in situations where it is meant to be used and where it can perform well. Figure 2 illustrates the relationship between overtrust, distrust and appropriate trust.

Based on the trust in automation literature, we see that operator trust in automation is affected, among others, for example the following factors:

- Operator's conception about the capabilities of the automation
- Operator's conception about the reliability of the automation (e.g., error-proneness according to previous experience)
- Operator's familiarity with the system
- Operator's self-confidence and trust in his or her own skills

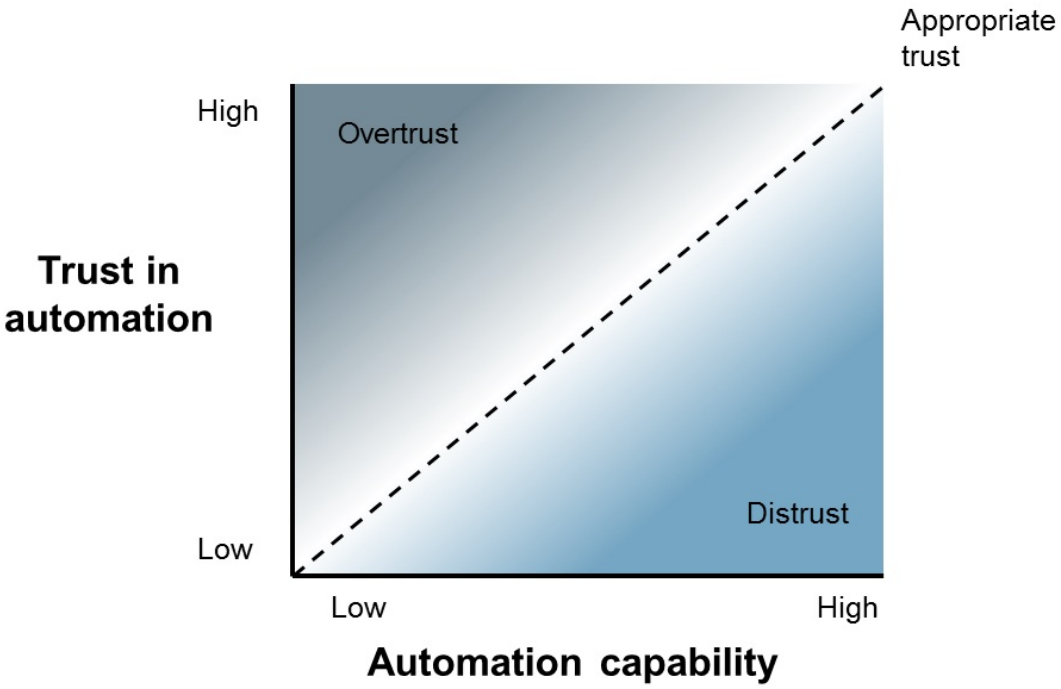

Figure 2. Appropriate trust

- Operator’s current state of vigilance, fatigue and workload
- Operator’s cultural and organizational background
- Feedback given by the automation about the functioning of the system
- Predictability of the system
- The complexity of the task
- Usability, user experience and user interface design of the system
- System delays
- System’s operating reliability due to high level of system transparency

To support the building the approppropriate level of trust in computer systems, various issues have to be considered in the design of the technology. The following are some design guidelines found from the literature regarding trust in automation (Lee and See, 2004):

- Show the system’s past performance
- Show the systems’ processes and algorithms in an understandable format to the operator
- Simplify the system’s algorithms and functioning so that it will be more understandable for the operator
- Show the automation’s purpose, design basis, and range of applications in such a way that they are related the operators’ goals
- Train operators about the functioning mechanisms (governing the behaviour), intended use, and expected reliability of the system
- Consider very carefully whether to make the system anthropomorphic

- Train the operators about automation's capabilities in different situations

4.2 Cryptography for Human Senses

Traditionally, secure cryptography is something that requires a great deal of computational effort and is thus only suited for machines and computers to process. Some advances have been made in providing users ways to perform and check connections between devices through visual comparison of numeric values or by bumping devices together. However, these have very limited use cases and are thus not general purpose methods.

Visual cryptography [18] has shown promise in giving users the power to decrypt encrypted data merely by looking at images. This has been expanded in other work, but it is not yet utilised in many applications. There are also methods for visual hashing[1], where one way hash function output is visualised to help users compare results. This has been used for example in the n-Auth mobile authentication scheme [20].

Other senses have not been used in conjunction with cryptography. Some senses might be difficult to imagine having cryptography translated to suit them. We think that at least auditory and touch (or haptic) senses could be utilised in similar manner as vision with visual cryptography. For example, one could have Braille writing that is formed from two different sources of "dots" that in themselves are not meaningful, but in combination form readable Braille characters like with visual cryptography. Also sound waves and superposition might offer similar possibilities.

5. Our Vision

The vision that we set out to target with our project was to build methods to realise HVC. To this end, we have some demonstrators that show some of the key technologies that we used. In addition to this, we have begun to formulate a protocol for HVC. This work is still ongoing, and will require more collaboration and research on many levels.

Our demonstrators show the envisioned functionality of some technologies that could be applied in HVC. Our grand vision is, that a user would have in every situation some method for verifying the computations made in AR/VR systems via human senses. The verification could be done via visual, audio, haptic or even other sensory methods that are available in the system.

We also need to have protocols that bridge the gap from the machines to human users in a secure manner. In HVC we need to include the human user in our models. Thus the traditional *client-server* modeling of systems is not necessarily the best, as the "client" in our end is not a similar machine as the "server" or the larger system that the user interacts with. Thus, models that can incorporate human users, with their strengths and weaknesses are needed. In our opinion, none of the current protocols take this into account, with the minor exception of PRISM [4], where the protocol and system is designed keeping in mind human limitations.

In Figure 3 there is a high-level overview on the different interacting parties in an AR/VR system. It also shows how the trust could be built and represented to the users in such an environment. It is noteworthy that the actual service and the rendering of the AR and VR are separated and this could provide tools in realising our vision. There are only few rendering engines at the

[1] for example `https://github.com/thevash/vash`

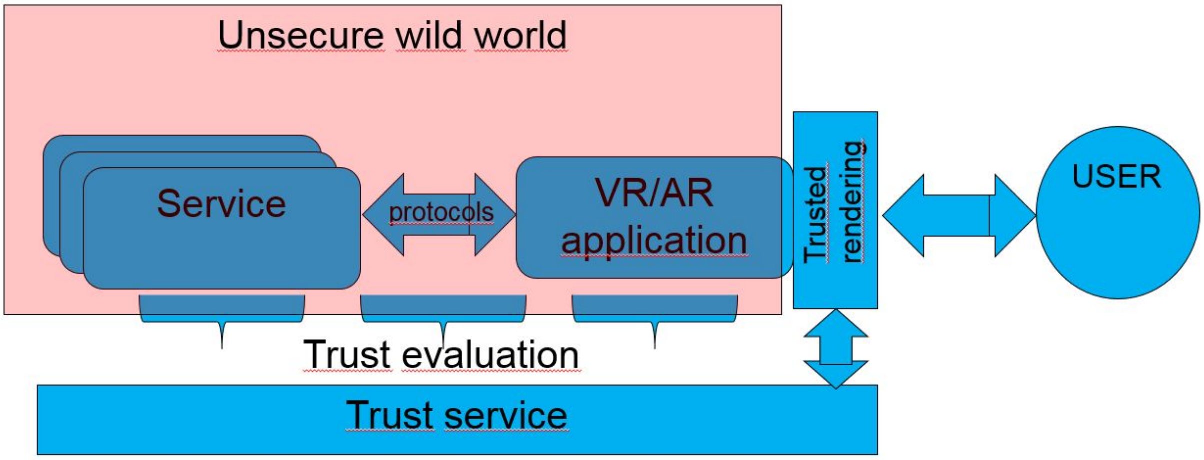

Figure 3. A high-level architecture for HVC in AR and VR

moment and these could be secured better than all the myriad services that can be located in the AR and VR environments.

6. Future Research

Even with the progress that our project has shown in this area, the question is still far from being solved. In the future at least the following questions should be solved in order to move towards a comprehensive Human Verifiable Computing solution.

- What senses are best to utilise in conveying security and trust information to the users?
- How can we ensure that the user is responsive to the information? Are there contextual cues that can be used to adapt the method(s) of dissemination?
- How to realise "visual cryptography" for other senses as well?
- What are the minimal trust assumptions that we can achieve for Human Verifiable Computing? Are these usable? What trust assumptions are easiest to work with?
- How to revoke trust in an intuitive and easy way across different services and platforms?
- How to formulate protocols for HVC? How to verify that these protocols are secure?

The above list is by no means exhaustive and we believe that there will be other topics and research questions will emerge, when this problem is being tackled by different parties.

It is noticeable that many of the technical subareas are still topics for ongoing or future standardization activities. MPEG (ISO/IEC JTC1 SC29) seems to be the most important forum in the area. Note that ISO/IEC JTC1 is also active in developing privacy, security and encryption related standards. Thus HVC research could have an impact in these developments.

We hope that we can begin solving these research questions in collaboration with companies, researchers and other interested parties. With the help of HVC, we can bring more trust, the most essential ingredient in effective communication, to human machine interactions.

References

[1] A. Aijaz, M. Dohler, A. H. Aghvami, V. Friderikos, and M. Frodigh. Realizing the tactile internet: Haptic communications over next generation 5g cellular networks. *IEEE Wireless Communications*, 24(2):82–89, 2017.

[2] N. Bostrom and E. Yudkowsky. The ethics of artificial intelligence. *The Cambridge handbook of artificial intelligence*, pages 316–334, 2014.

[3] A. Cavoukian and J. Jonas. *Privacy by design in the age of big data*. Information and Privacy Commissioner of Ontario, Canada, 2012.

[4] J. Franklin, M. Luk, A. Seshadri, and A. Perrig. Prism: enabling personal verification of code integrity, untampered execution, and trusted i/o on legacy systems or human-verifiable code execution. *CyLab*, page 41, 2007.

[5] K. Halunen, J. Häikiö, and V. Vallivaara. Evaluation of user authentication methods in the gadget-free world. *Pervasive and Mobile Computing*, 2017.

[6] B. Hannaford and A. M. Okamura. Haptics. In *Springer Handbook of Robotics*, pages 1063–1084. Springer, 2016.

[7] V. Hayward, O. R. Astley, M. Cruz-Hernandez, D. Grant, and G. Robles-De-La-Torre. Haptic interfaces and devices. *Sensor Review*, 24(1):16–29, 2004.

[8] O. Jason, K. Kiyoshi, and T. Haruo. Virtual and augmented reality on the 5g highway. *Journal of Information Processing*, 58(2), 2017.

[9] M. Kosinski, D. Stillwell, and T. Graepel. Private traits and attributes are predictable from digital records of human behavior. *Proceedings of the National Academy of Sciences*, 110(15):5802–5805, 2013.

[10] K. Lebeck, T. Kohno, and F. Roesner. How to safely augment reality: Challenges and directions. In *Proceedings of the 17th International Workshop on Mobile Computing Systems and Applications*, pages 45–50. ACM, 2016.

[11] J. D. Lee and N. Moray. Trust, self-confidence, and operators' adaptation to automation. *International journal of human-computer studies*, 40(1):153–184, 1994.

[12] J. D. Lee and K. A. See. Trust in automation: Designing for appropriate reliance. *Human factors*, 46(1):50–80, 2004.

[13] S. Lewandowsky, M. Mundy, and G. Tan. The dynamics of trust: Comparing humans to automation. *Journal of Experimental Psychology: Applied*, 6(2):104, 2000.

[14] J. M. McCune, A. Perrig, A. Seshadri, and L. van Doorn. Turtles all the way down: Research challenges in user-based attestation. Technical report, DTIC Document, 2007.

[15] P. Milgram and F. Kishino. A taxonomy of mixed reality visual displays. *IEICE TRANSACTIONS on Information and Systems*, 77(12):1321–1329, 1994.

[16] B. M. Muir. Trust between humans and machines, and the design of decision aids. *International Journal of Man-Machine Studies*, 27(5-6):527–539, 1987.

[17] B. M. Muir and N. Moray. Trust in automation. part ii. experimental studies of trust and human intervention in a process control simulation. *Ergonomics*, 39(3):429–460, 1996.

[18] M. Naor and A. Shamir. Visual cryptography. In *Advances in Cryptology—EUROCRYPT'94*, pages 1–12. Springer, 1995.

[19] M. A. Otaduy, A. Okamura, and S. Subramanian. Haptic technologies for direct touch in virtual reality. In *ACM SIGGRAPH 2016 Courses*, page 13. ACM, 2016.

[20] R. Peeters, J. Hermans, P. Maene, K. Halunen, K. Grenman, and J. Häikiö. n-auth: Mobile authentication done right, 2017. Accepted to ACSAC 2017 conference.

[21] T. Pfister, X. Li, G. Zhao, and M. Pietikäinen. Recognising spontaneous facial micro-expressions. In *Computer Vision (ICCV), 2011 IEEE International Conference on*, pages 1449–1456. IEEE, 2011.

[22] F. Roesner, T. Denning, B. C. Newell, T. Kohno, and R. Calo. Augmented reality: hard problems of law and policy. In *Proceedings of the 2014 ACM international joint conference on pervasive and ubiquitous computing: adjunct publication*, pages 1283–1288. ACM, 2014.

[23] N. Sebe. Multimodal interfaces: Challenges and perspectives. *Journal of Ambient Intelligence and smart environments*, 1(1):23–30, 2009.

[24] J. Song, J. H. Lim, and M. H. Yun. Finding the latent semantics of haptic interaction research: A systematic literature review of haptic interaction using content analysis and network analysis. *Human Factors and Ergonomics in Manufacturing & Service Industries*, 26(5):577–594, 2016.

[25] A. Tiwari. A review on haptic science technology and its applications. *Science [ETEBMS-2016]*, 5:6, 2016.

[26] K. J. Vicente. *Cognitive work analysis: Toward safe, productive, and healthy computer-based work*. CRC Press, 1999.

[27] M. Walfish and A. J. Blumberg. Verifying computations without reexecuting them. *Communications of the ACM*, 58(2):74–84, 2015.

[28] S. Zuboff. *In the age of the smart machine: The future of work and power*. Basic books, 1988.