

Reflections on Trust in Devices: An Informal Survey of Human Trust in an Internet-of-Things Context

Geir M. Kjøien

Published online: 23 August 2011
© Springer Science+Business Media, LLC. 2011

Abstract Trust is an essential component of every interaction we make. For some transactions we need very little trust, while for others we may be quite risk averse. In this paper we investigate trust in an Internet-of-Things environment. Can we trust devices? How can we quantify trust in devices?

Keywords Trust · Internet-of-things · Software · Hardware · Subjective logic

1 Introduction

In this paper we examine aspects and issues related to trust in devices in an Internet-of-Things (IoT) environment. The primary focus is on human trust in devices.

1.1 The Internet-of-Things

The IoT environment is a diversified and heterogenous environment, in which there will be a multitude of different devices. Some of these devices will be personal devices like our mobile phones and some will be anonymous and expendable devices, some will be high-assurance and others will be cheap and unreliable, some will enjoy physical protection while others will be distributed into hostile environments. What they all have in common is (a) the ability to communicate and (b) the ability to compute results. So the IoT devices will be small connected computers, some highly dedicated and others more general purpose. They will also have in common that they perform some service.

1.1.1 Assumptions

In this paper we primarily investigate trust aspects of Human-to-Machine interactions. Many of the observations may apply to Machine-to-Machine as well. As a generalization we

G. M. Kjøien (✉)
University of Agder, Kristiansand, Norway
e-mail: geir.koien@uia.no

assume that device access is wireless (i.e. not directly observable to a human consumer), that devices are ubiquitous and that their identity or address is unknown to the human consumer. We therefore have that the human consumer in general cannot be 100% certain with which device he/she is interacting and furthermore that she/he does not know the identity and/or address of the device. We may assume that the human user is aware of the service being requested and we may also assume that the human user has indication about what type of device he/she is interacting with (or intended to be interacting with).

1.2 Trust and Trustworthiness

1.2.1 Trust

By trust we mean reliance on the integrity, ability or character of an entity. Control and necessity plays a role here, and we may be compelled to rely on entities due their imposing authority (control) or to lack of alternatives (necessity). Trust can be further explained in terms of confidence in the truth or worth of an entity. The catch here is that we need to apply this to human assessment of “truth or worth” of distinctly non-human devices.

1.2.2 Trustworthiness

This concept is related to trust and points to worthiness of trust or belief. An entity being trustworthy implies that other parties believe that the entity will take responsibility for its actions, its conduct and its obligations. There must of course also be willingness and ability to comply with this belief on part of the trustworthy entity. Needless to say, the trustworthiness is relational and relative, and it may mean different things to different entities.

For humans there is an ethical dimension to being trustworthy; this notion clearly does not apply to devices. Instead, for devices, one is left with a probabilistic notion of intention and ability (i.e. the probability that the device will “intend” do as expected combined with its ability to actually do so). The “ability” to behave as expected must be present in under adverse conditions.

1.3 Outline of this Paper

In the next section (Sect. 2) we briefly investigate limits to trust in software, hardware and devices in general. Then in Sect. 3 we analyze the IoT environment in more detail, before we in Sect. 4 informally reflect on human trust. Human trust is hard to quantify, but in Sect. 5 we investigate a promising method for doing just that. The method, which is not limited to human trust in devices, is based on subjective logic. It cannot alone solve the problem of trust for us, it may at least provide us with a tool to enhance our understanding and to enable coarse grained estimations of trust. We conclude the paper in Sect. 6 with a brief summary, possible directions for further research and a few concluding remarks.

2 Reflections on Trust in a Digital World

A working IoT device consists of physical hardware (processor, memory, I/O hardware, sensors/actuators etc), software (firmware, operating system, drivers, applications) and last but not least a power source. Trust in an IoT device implies that these components themselves must be reliable and trustworthy, and comply at least to some minimal standard of operation.

The device and its components must not only act as expected, but they must be able to do so in an hostile environment. In this section we investigate trust and trust worthiness with respect to the software and hardware components. We also briefly examine the mobile device, which may form the basis for many IoT devices, and the trusted platform concept.

2.1 Limits of Trust in Software

2.1.1 *Reflections on Trusting Trust in Software*

In his 1984 ACM Turing Award acceptance speech Ken Thompson presented the paper “Reflections on Trusting Trust” [1]. Thompson outlines a scheme in which an intruder has access to the source code of a C compiler and then modifies the C compiler to contain Trojan code. The attack code is such that when the compiler compiles the `login` program it will make `login` accept either the intended password or a predefined (by the intruder) password. This is in itself a clever little hack, but as noted by Thompson:

Such blatant code would not go undetected for long. Even the most casual perusal of the source of the C compiler would raise suspicions.

However, the real trick is that there is a second piece of Trojan code. This code is aimed at the C compiler itself. The “clean” original C compiler (binary) is used to compile the infected compiler source and produces a Trojan (binary) C compiler. Then, the original C compiler source is restored, but whenever anyone tries to recompile the C compiler, the infected binary will see to it that the new resulting C compiler binary is equally infected. Then, no amount of source code inspection will ever reveal the Trojan.

So what can we learn from the above little trick? According to Thompson:

The moral is obvious. You can’t trust code that you did not totally create yourself. (Especially code from companies that employ people like me.) No amount of source-level verification or scrutiny will protect you from using untrusted code.

The above tale emphasizes two important aspects with respect to trust and computer programs, namely that (a) One cannot trust programs (binaries) and (b) that one cannot trust programs with source code either!

2.1.2 *What About Open Source?*

With the C compiler example in mind one may ask: To what extent does it matter whether or not the adversary has access to the source code? And perhaps; Is the problem intrinsic to open source software? The fact that the source code was available appeared to be part of the problem, but of course it could also be part of the solution (“to many eyes, all bugs are shallow”[2]). In the literature one will find arguments both for and against “open source” with respect to security, and indirectly to the trustworthiness. Anderson, using reliability growth theory, has found that there is no definitive answer [3], and that the pros and cons tend to cancel each other out. So, it seems that open source is neither the problem nor the solution with respect to security flaws, and by association we have no reason to trust open source software more or less based on this one property alone.

2.2 Limits to Trust in Hardware

So, we know that we cannot entirely trust software. The solution, it seems, would then be to have dedicated hardware that enforce security and security policies. This type of hardware exists and provided that the secure environment hardware indeed performs as expected, many software security issues may be solved or mitigated. However, even putting aside philosophical arguments about decidability and executability, one should realize that hardware cannot guarantee security. Hardware is not immune to attacks and with respect to trust in hardware one should bear this in mind.

2.2.1 *The Trusted Platform Module and Other Trusted Environments*

One attempt to provide a secure execution environment is the so-called Trusted Platform Module (TPM). The TPM concept is standardized and it details a secure cryptoprocessor, which has secure storage and provides secure processing. The TPM concept is specified by the Trusted Computing Group and transposed to an ISO/IEC standard (ISO/IEC 11889 [4]).

The smart cards used in mobile stations in the 3GPP cellular systems (SIM card/UICC card) is another example of a trusted execution platforms [5]. The 3GPP has also defined a trusted environment (TrE) to be used in 3GPP-based femto-cells [6].

2.2.2 *Enter the Hardware Trojan*

The processing unit in a typical IoT device need not be the most powerful, but it may nevertheless contain millions of transistors. This brings up the question about how can one be assured that the circuits doesn't contain malicious instructions?

In the paper "Hardware Trojan: Threats and Emerging Solutions" [7] the authors outlines the concept of a hardware Trojan, and with it means to assurance and trust in safety-critical applications. They point out that automated design tools could insert malicious instructions, much like the instrumented C compiler in Thompson's example that corrupted the `login` program. They also point out that new designs are typically built around pre-made building blocks with automated wirings etc in place. Thus, unless one is fully assured of all these tools (software and hardware) and the fabrication process itself, we have that (to paraphrase Thompson) "You can't trust hardware that you did not totally create yourself".

In [8] the authors presents a taxonomy of hardware Trojans and how they might be implemented and in [9] the authors report their experiences in designing and implementing hardware Trojans. So, while hardware Trojans may be a bit exotic and while they perhaps are a somewhat remote threat today, in the future one may find hardware Trojans being used in cyber warfare or in cyber crime like the infamous Stuxnet virus/Trojan [10].

All in all we have that hardware Trojans are feasible or becoming feasible, and that they effectively impose a limit to the trust we may have in hardware.

2.3 Trust in Personal Devices

We have now established that one cannot have full trust in software, which after all was no surprise. We have additionally demonstrated that malicious code can exists in hardware too. Thus, even a trusted platform module cannot be fully trusted as-is. With this as an upper ceiling to our trust, what then can be said of personal devices?

Table 1 Scope of link layer protection in 3GPP-based systems

Protection\system	2G/GSM	2G/GPRS	3G/UMTS	4G/LTE
Range	MS-BTS	MS-SGSN	MS-RNC	MS-eNB
Confidentiality	64-bit	64-bit	128-bit	128-bit
Integrity	n/a	n/a	128-bit	128-bit

2.3.1 The Mobile Phone

Most mobile phones today are relatively simple devices, but this is changing fast, and the smartphone is becoming the dominant platform in mature markets. What signifies a smartphone? Well, apart from providing seamless cellular connectivity and more, a typical smartphone is a reasonably powerful computing platform, it runs an advanced operating system, it has plenty of memory and it has quite a few applications. The applications may come from third parties and they may be downloaded on demand by the consumer.

2.3.2 Trust in the Mobile Phone Communications Security

There is a common misconception that mobile phones provide end-to-end security. They do not, and in fact the original design rationale for providing cellular/wireless security has tended to be to provide over-the-air protection only. E.g., the 3GPP systems basically provide the users with (over-the-air) link layer security. Table 1 provides an overview of range and coverage for the 3GPP systems.

Note that Table 1 is slightly misleading in that the 3G/4G data integrity protection is only for the control plane. Furthermore, while the keys are 128-bit wide the integrity checksum is only 32-bit wide. Access security in the 3GPP systems is presented and analyzed in [11]. Overall we find that the protection is limited to the wireless segment in the access network.

2.3.3 Trust in the Mobile Phone Device

As with any hardware device the mobile phone may be compromised by a hardware Trojan, but while this is possible we do not see it as a very realistic scenario today. Thus, the mobile phone hardware can likely be trusted, but one should bear in mind that the hardware will contain test circuitry, digital rights management circuitry etc that may be used against the user by an intruder.

2.3.4 Trust in the Mobile Phone Software

There is no reason to assume that smartphone software is any worse or better than software for other general purpose computing platforms. This assumption is supported by the fact that the Apple iPhone OS software is partly modeled on Apple's MAC operating systems, that the Android OS is based on Linux and the Windows Phone 7 OS is based on the Windows 7 platform. The applications running on these platforms are tempting targets for hackers, and so it seems prudent to assume that they are about as vulnerable to malicious software as would be generic personal computing software.

2.3.5 Trust in the Secure Modules (Smart Cards)

The SIM card or the more recent UICC [5] represents a trusted environment for the mobile phone. However, its use is mainly limited to security functions for cellular link protection [11]. The UICC may have additional functionality to facilitate banking and eCommerce. However, one must be cautious to trust UICC-based functionality too much, since all user I/O to the UICC is by means of the mobile phone. Furthermore, there is generally no authentication between the UICC and the mobile phone, and so one cannot trust UICC-based application that rely on user I/O unless one trust the mobile phone too. This does not mean that UICC-based security is without merit, it simply means that trust in UICC-based security must be conditional. We add to this that trust in the UICC as a platform implies trust in the UICC manufacturer as well as in the mobile network operator (which issues the UICC).

2.4 Trust in Security Procedures and Software Updating

2.4.1 Trusting Security Updates and Updates of System Functions

Can the security procedures/products be trusted and are they trustworthy?

It may seem that the security procedures/products *must* be trusted and that they indeed should be trustworthy. However, security is only as static as the environment and it is an everyday experience that there are “security updates” and “virus definitions” to be installed. This of course serves to highlight the fact that security functionality is (a) not perfect and error free and that (b) security is a moving target and as such can only be trustworthy if it is updated regularly.

To update security procedures or for that matter other core system functions is clearly itself an operation that needs to be secured and fully trustworthy. When it comes to basic security functionality provided by the operating system we can be relatively confident that there is a formal procedure in place. Furthermore, we may also assume that the security updating functions are themselves secured and that they (functionally) have been tested and will work as expected. Thus, for system updating functionality we may expect the intentions to be good and the ability to be strong. But not perfect.

2.4.2 Updating Applications

For dedicated IoT devices we expect there only to be a few selected applications running on the device. On the other hand, the typical smart phone may support a plethora of applications (apps). Depending on the platform, the apps may be updated via a platform function (for instance via Android Market), but ultimately the quality of the update will depend upon the developers.

2.4.3 Trust in Updating Procedures

There will in principle be two types of updates that can be performed; (a) Core function updates and error correction (including security) and (b) Functional updates (providing new functionality).

Core function updates will often be seen as overhead and unless they are critical, and there may be a tendency to delay updating, possibly also to ignore updates. However, core

updates may also be so important that they are made mandatory. Functional updates may bring their own incentive, and as such it may be more likely that they are installed in a timely fashion. From a security point of view we observe that attackers may exploit this, and there is no shortage of fake updates that tries to trick human users to download and install malware.

Trustworthiness, with respect to the updating procedures, will depend (as a minimum) on the following:

- Availability/timeliness of updates
- Correctness and completeness of updates
- Ability to securely download and apply updates

For devices with limited bandwidth, limited computing power and limited power (battery provided) the last item is important. The updating cost (processing/bandwidth/power) may prohibit frequent updates and it may under many circumstances be infeasible to download and apply the updates.

2.5 Trust in IoT Devices and Services

2.5.1 *Trust Control in an IoT Environment*

The IoT devices will be plentiful and diverse, but as a common minimum they will all communicate and they will all be able to process and respond to requests. By default we may assume that the IoT devices have similar characteristics to our personal device discussed in the previous subsection (Sect. 2.3). Thus, we know that people tend to trust the device, but that the trust may not always be warranted. The “personal” aspect of a personal device may mean that people protect the device and that they for better or worse trust it more than they would with an impersonal device.

Thus, what we have is that the physical properties and the security properties may be the same or similar for personal and impersonal devices alike, but the perception about their trustworthiness may vary. Similarly, the threats, the vulnerabilities and the exposure level may vary significantly.

The IoT environment itself will be discussed in more detail in Sect. 3, but as noted by Liu and Wang [12], trust control in heterogenous network environments like the IoT environment is important, and furthermore it seems important and evident that trust control must be supported such that it is not left to the human user alone to make the decisions.

2.5.2 *Trust in Services*

In [13] the authors report from a survey of trust in internet applications and internet services. This may not be directly comparable to IoT devices and the associated services, but we believe the match is sufficiently close to warrant our interest. So what did the authors find out? Well, they classified trust, they examined security policies, and they investigated issues and aspects of authorization and security management. They found that trust relationships are particularly important and that trust management systems is a key element. Amongst the findings are that systems that employ credible entity authentication and access control are also the system most likely to be trusted by the users/consumers. That is; Systems demonstrating ability and willingness to protect themselves are the same systems that one would tend to trust.

3 The Internet-of-Things Environment

3.1 Dynamic Heterogeneous Environments

New devices will be deployed continuously and through upgrades old devices will be able to perform new services. Similarly, the IoT environment will be adaptive to consumer needs and it will likely also be adaptive towards threats in the environment. For example, security measures such as secure routing, firewall rules, protection against malicious software and similar will be updated and upgraded according to the present “risk climate”. Furthermore, the heterogenous nature of both the status of the devices (software revision, security updates etc) will make it extremely hard defend the IoT devices adequately at all times, but it will also make it very hard for an intruder to successfully attack all devices. This is evidenced in attacks with malicious software, which often depends crucially on the state of the software (OS and applications), the status of security defenses (security updates, anti-malware software etc).

3.2 Hostile Environments and Exposure Control

3.2.1 *Hostile Environment*

The IoT environments will be diverse, and some will be very hostile. Still, some (if not most) of the environments will be quite safe and well protected. The real problem is that the user or the IoT device owner/operator cannot necessarily know which is which, and that the same environment may be hostile to one party while it is reasonable safe to another party. A point in case may be the so-called home-cell/femto-cell concept, in which a cellular basestation for UMTS/LTE can be deployed at the customer premise. To the customer the deployment will be in an environment in which he/she potentially have full control. On the other hand, dishonest subscribers exists and they may attempt to intrude on the device. To the operator the consumer premise may therefore potentially be a hostile environment.

3.2.2 *Exposure Control*

It would seem good advice that honest parties should try to avoid hostile environments, but this is easier said than done. It is not clear that we are indeed able to recognize a hostile environment. Furthermore, given the dynamic nature of the IoT environments, the threat exposure may vary considerably over time. Thus, we need to be able to dynamically adjust our trust and our defences. To this end it important to control the exposure level. Exposure control therefore implies continuous monitoring, dynamic risk assessment and adaptive mitigation strategies.

3.3 Risks Analysis

Risk analysis is often provided as a snapshot of the risk level for a given system at a given time. As such it may not be too useful to once-and-for-all assess risks and threats for an IoT environment, since we expect it to be highly dynamic.

Anyway, one would start off with identifying the assets to be protected and with identifying the stakeholders. Then, one estimates the value of the assets, realizing that the value would be relative to each stakeholder. With this in mind one then proceeds to identify potential risks and threats towards the assets, and if possible assign likelihood and estimate potential

impact. In this respect there is nothing special per se to risk analysis for IoT devices and IoT services.

3.4 Intruder Models for an IoT Context

There exists several theoretical intruder models, amongst them the classical Dolev-Yao Intruder (DYI) model [14]. The DYI is a very powerful intruder which has immense computational power and foresight to use all available knowledge to intrude on the system. However, the DYI is incapable of physical intrusion and it is unable to break cryptographic primitives. In real life cryptographic primitives may be broken and physical intrusion is a likely risk prospect. We advice aiming at full protection against logical (protocol level) attacks. Thus we advice retaining the DY Intruder model for the IoT context, but also advice system designers to be vary of the DY Intruder model limitations. Detection and mitigation will be the order of the day.

3.5 Security Measures

Security measures for IoT devices, IoT services and the consumers should reflect the risks and threat one has identified. Intruders should be taken into account and one should perform cost-benefit analysis of protective measures to be deployed. The cost-benefit analysis should take a long-term view and should allow for some margin with respect to the protective measures. The mainstay of computer and communications security methods, entity authentication, data confidentiality and data integrity, is still the most likely to be deployed. Depending on the device, other protection measures such as anti-malware and firewalls may also be deployed. As indicated in Sect. 2.5 deployment of suitable security protection methods is important for consumer trust. We argue that use of hardware-based trusted environments etc (see Sect. 2.2.1), while not perfect, are necessary and that reliance on software-only protection is untenable in open environments.

Nothing very special with IoT here at all, and security for IoT systems would likely benefit from complying with principles and methods used in cellular systems or similar.

4 Aspects of Human Trust

The following is an informal and non-exhaustive investigation of human trust. It will serve to identify and highlight some central aspects of human trust.

4.1 Transitivity and Reflexivity

Transitivity and reflexivity with respect to trust is about how our trust affects and are affected by our surroundings. All trust-related transactions we have with other entities and objects automatically assesses these aspects of trust, whether we are consciously aware of it or not.

4.1.1 Trust and Transitivity

Trust is clearly transitive to some extent, and this is commonly used in actual systems. For instance, in cellular communications we have that a subscriber (S) trusts the home network (H). The home network (H) in turn has some level of trust in a visited network (V). The

subscriber S can therefore afford some level of trust in the visited network V , but this trust is indirect and it should under no circumstance be (initially) stronger than the weakest link.

4.1.2 Trust and Reflexivity

Trust may not necessarily be reflexive. The fact that A trusts B does not imply that B should trust A .

At least not initially. But, with human relationships there is often a tendency of mutuality or equality. There is also social pressure at work here, and it is often understood or required that trust should be more or less reflexive amongst peers. The peer concept may here be interpreted as someone/something being more or less equal in power. In security parlance, one may say that the parties should have more or less equal control (or jurisdiction). As the asymmetry in control, power or authority increases one should be less and less inclined to assume that the trust is reflexive.

4.2 Psychological Aspects of Risk and Risk Assessment

To recognize risks and to assess them realistically is not easy. However, humans are actually well adapted to recognizing and relating to some types of risk. These would be risks we would historically have encounter when man lived in tribes or as small family groups. Risk recognition and risk assessment in a digital age is something quite different, and our brains seems to play tricks on us in this respect. In the papers [15, 16] the authors analyze the psychology of security, or more specifically our ability to recognize, differentiate and assess risk. As it turns out we have problems analyzing abstract risks like those associated with computers and the internet, we tend to believe that bad thing happens to other people and not us (wishful thinking) and we are bad at assessing severity. As reported by West [15]:

People tend to believe they are less vulnerable to risks than others. People also believe they are less likely to be harmed by consumer products compared to others. It stands to reason that any computer user has the preset belief that they are at less risk of a computer vulnerability than others.

Schneier report similar attitudes in [16], and also points to our inability to deal rationally with spectacular but unlikely events vs. mundane and likely events:

- People exaggerate spectacular but rare risks and downplay common risks.
- People have trouble estimating risks for anything not exactly like their normal situation.
- Personified risks are perceived to be greater than anonymous risks.
- People underestimate risks they willingly take and overestimate risks in situations they can't control.
- Last, people overestimate risks that are being talked about and remain an object of public scrutiny.

In the context of a human user and an IoT device, we note that the above may predispose humans to trust commonly used devices and services, in particular if the encounter represents an everyday event. Devices and services we seldom use, or for cases that are perceived to be outside our control, may be distrusted since we assume a too high risk. For both cases there may be a mismatch between the actual risk level and the human trust in the device.

4.3 Distrust, Deception, Retaliation and Altruism

4.3.1 *Distrust*

In addition to have trust in someone or something, one may also lack trust. So if we express trust numerically, we may have positive values and we may have neutrality at zero. Then, it follows, it is also possible to have negative trust or distrust. It is not clear whether or not we would necessarily behave differently towards a person or thing that we distrust or simply lack trust in, but we might very well influence our surroundings (our trust network) differently for these cases. Indifference is one thing, but active distrust quite another matter.

4.3.2 *Deception and Retaliation*

Conscious deception is an act commonly attributed to humans and the human brain. Psychologists often refer to this in contexts such that we have a theory of mind, and that this theory of mind allow us to construct sophisticated attacks against other people based on deceptions [17]. Of course, an advanced theory of mind also allow us to protect ourselves against the very same attacks. In an evolutionary context deception and “parasitic” behavior may turn out to be a viable strategy under some circumstances (the cuckoo example). However, humans tend to recognize and remember other people and objects they encounter, and the reciprocal principle then comes into play (tit-for-tat rule). Dawkins discusses these quite complex phenomena in detail in [18, 19].

The pay-off of deception may be immediate, but the cost of retaliation will eventually come. However, here we have the preconditions that (a) the deceived party must realize that he/she has been deceived, (b) the deceived party must recognize and remember who/what was behind the deception, and (c) there must exist an opportunity for revenge or retaliation.

Obviously, “misdemeanor” type of offenses may not warrant much retaliation and the human willingness to conduct retaliatory behavior often tend to diminish over time. Of course, humans are sociable beings and we are strongly influence by our culture and our friends etc. A feeling of disrespect and loss of face may be associated with deceptive behavior, and so a deception may therefore have a symbolic value far exceeding any economic losses, and the willingness to retaliate and retaliate strongly is influenced by this. There may therefore be rather complex prisoner’s dilemma type of difficulties involved in deceptive behavior and our reactions to it.

The above notion of deception applies primarily to human-to-human interactions, but as we have a phycological tendency towards personification of inanimate objects it may also, at least in part, apply to human-to-device interactions. However, apart from physical destruction etc, how might a human retaliate in any meaningful way towards a device or service?

Well, we might expect a human which feels deceived by a device/service to trust the device/service less in the future. Loss of trust may lead to active distrust and retaliatory behavior, possibly in the direction of attacking the reputation of the device/service (and by association to the service provider/device operator). This is in fact quite likely, and one often sees examples of this in consumer groups on various social networks on the internet.

4.3.3 *Altruism*

Altruism seems to be inherent in human beings, but with qualifiers and preconditions. Again one may find evolutionary explanation for altruism [19]. With respect to trust this disposition towards altruism would mean that a human neutral stance on trust may in fact be skewed

slightly towards positive trust. In nature we find reciprocal altruism, and as with revenge and retaliation, there are conditions such as the ability to recognize altruistic behavior, to attribute it to an entity or object and to remember it.

Sophisticated scams and deceptions may actively manipulate us by invoking reciprocal altruistic behaviors. The intruder could exploit this by first building trust with initially honest behavior, before exploiting the trust in subsequent interactions. Exploitations of reciprocal altruism may also be more mundane and less offensive, e.g. as part of a PR strategy or a marketing effort.

4.4 Reputations, Association and Brands

It should not come as a surprise that reputation matters. The reputation could be good or bad, it may be well founded or not, it may be fixed or circumstantial, but reputation is clearly a decisive factor for trust.

Association also matters in this respect. That is, new services or unknown devices associated with something or someone we know will gain or lose from that association. Thus, we may assume that there is a certain level of transitivity and reflexivity with respect to reputation. This is of course used extensively in marketing, where one commonly attempt to associated new products with trusted or admired brands or products. The association part also works in that many tend to trust what other people say, in particular if they are famous or known experts. The association with “desirability” also comes into play in that we are influenced by what we desire (Ref. lifestyle advertisements). In the paper “Trust and the Internet of Things” [20], the authors examine the reputation concept by allowing people to express opinions by tagging in a virtual environment. Trust is here a precondition for beliefs in other peoples opinions, and trust is indeed also affected by other people and their opinions.

4.5 The Human Brain

Obviously, we will not go into much detail about the human brain in this paper, but a brief overview may provide some pointers as to why our decisions regarding security, risk and trust so often seems based more on emotional responses than rational thinking. For exhaustive treatment of the human brain, cognition and emotional responses we refer to [21].

Human beings are capable of abstract thought, but much of our mental machinery is still directly geared towards physical interactions with the world and emotional responses to stimuli. Amygdala is an “ancient” part of our brain close to the brain stem, and it deals amongst others with emotional responses (fight-or-flight decisions etc) and it affects social and emotional behavior. The responses are reflexes and more or less “unconscious”.

The Neocortex is the newest part of the brain, and it is directly associated with consciousness, reasoning, thinking and abstract thought. However, compared to Amygdala the Neocortex is fairly slow in its decision making. In the modern world quick responses are still valuable, but decision making concerning new, abstract and hard to assess digital threats and risks are often best left to the Neocortex.

Many scams rely on manipulating our emotional responses. This is sometimes referred to as the “Amygdala hijack” [22]. The Neocortex can overrule decisions by the Amygdala, but it takes an effort. So, unless the decision in the Neocortex is strong it will not overrule the Amygdala. Likewise, if we have been strongly emotionally affected, even quite conclusive reasoning may not alter our decisions. With respect to human trust in devices, it alludes to a situation where strong emotions will override knowledge. Thus, if emotionally conditioned we may trust or distrust a device far more or far less than objective knowledge would

support. And, interestingly, by itself more knowledge is not necessarily enough to sway our convictions.

5 Quantifying Trust

5.1 Requirements for Quantifying Trust

We have in the previous sections, in some detail, explained why we cannot always trust software, hardware, devices and/or services. We have furthermore seen that the IoT environment is dynamic, complex and potentially quite hostile. Adding to all this, we have briefly investigated aspects of human trust and seen amongst others that emotional responses may overturn rational thinking.

So, then it seems prudent to ask: Is it really possible to quantify trust and will such quantification be meaningful? In the following subsections we briefly investigate a promising model for quantifying trust. In this model we define trust relationships as we would normally do, but the relationships in a *trust network analysis - subjective logic* model are more complex than the trust relationship models we usually see in the literature. In order to model reality in a convincing way the implemented models may become quite complex. Still, we want the model to appear transparent, to be convincing and to be tractable in practical use (computationally and w.r.t. ascertaining validity).

5.2 Trust Relationships, Trust Networks and Beliefs About Trust

In the paper [23] the authors investigate and apply subjective logic in conjunction with trust networks [Trust Network Analysis-Subjective Logic (TNA-SL)]. The trust networks consist of transitive trust relationships between people, organizations and possibly virtual agents (software agents). In our IoT perspective all these entities would relate to each other and communicate over IP based internet connections. With the aid of subjective logic one formalize trust relationships and assign numerical values to reputation or as subjective trust measurements. The trust between parties within the community can then be computed by analyzing the trust paths linking the parties together.

In subjective logic trust measures are expressed as beliefs, and subjective logic is used to compute trust between arbitrary parties in the trust network. A specific belief is called an opinion. One has concepts such as belief, disbelief and uncertainty, and it is assuring that one also allows for negative trust scores. Subjective logic is furthermore capable of taking into account aging of beliefs. Thus opinions have real-time properties, and one has less confidence in an old opinion than in a recent opinion. This enables TNA-SL to capture dynamic environments, where beliefs and uncertainties change with time.

A criticism of TNA-SL has been that it reduces a complex trust graph into a set of series-parallel graphs, but in [24] the authors compares TNA-SL to optimal trust network analysis, and the indication is that these methods arrive at similar results. TNA-SL and similar methods shows great promise and may turn out to be useful tools when assessing trust in the disparate and ever-changing IoT environment. However, it is vital that the expressed “opinions” is in correspondence with what we perceive of as reality. The map must reflect the terrain so to speak. Therefore, as for any complex and learning method (in a Bayesian network way), we suspect that tuning will be crucial for its performance. Thus, similar to monitoring for other security aspects (malware, fraud, etc) we advise that the trust evaluation functions be monitored.

5.3 Difficulties in Determining Trust

The methods outlined represents a way of determining reputation/opinion based trust. However, the reliability of the the TNA-SL approach clearly depends on the number of known relationships. To continually exchange and reevaluate trust levels/beliefs might be infeasible (bandwidth restrictions etc). The trust network would tend to be for humans that evaluate trust in devices. This means that the computational and communications complexity limitations would generally apply to human-to-human channels, or rather to the human proxy devices (likely a smartphone) and their channels. The actual overhead will of course be closely correlated to the update frequency of the trust evaluations, and would furthermore be influenced by whether reputations are distributed by push (perhaps broadcast) or pull methods.

6 Summary and Concluding Remarks

6.1 Summary

We have investigated and presented a multifaceted view of trust in software, hardware, devices and services. The Internet-of-Things represents a highly diverse set of environments. The environments will range from static to highly dynamic, from simple to complex and from friendly to hostile.

We have also investigated some aspects of human trust and what may affect our trust. These are complex matters and the human psyche plays an important role here. It has not been our intention to analyze trust in depth with respect to evolutionary biology, neuroscience or psychology, but we found reason to briefly investigate some aspects of these sciences in order to understand how the human mind deal with security issues, with risks and with trust.

In Sect. 5 we briefly examined trust networks and how applying subjective logic may be a useful way to model trust in an IoT environment.

6.2 Directions

We believe that TNA-SL could be a useful modeling technique for human trust in IoT devices and services, but realize that there are aspects and issues that needs to be addressed before one can feel assured that TNA-SL, or related approaches, truly are good approximations for trust in an IoT environment. We therefore want to encourage further research into this area, and in particular we encourage efforts towards tuning of TNA-SL models to mimic real-world scenarios in a realistic and cost effective way. Here we would like to see addressed how TNA-SL models behave in environments where communications concerning opinions and reputations is restricted. That is, how does TNA-SL behave when parts of the graphs are out-of-date.

Adding to this, we would also like to see analysis of networks in which there are asymmetries in the propagation of opinions and where one may have broadcasting of reputations/opinions. We also want to inspire research in trust hierarchies in order to reduce the TNA-SL trust graph to be evaluated. Thus we suggest that trust proxies be evaluated for TNA-SL networks. It might also be wise to investigate risk (with respect to the required trust level) as a factor in determining completeness and frequency of the trust/belief evaluation. This, it is hoped, might make TNA-SL like approaches feasible and practical.

6.3 Concluding Remarks

Human trust in IoT devices and services is essential to the success and longevity of the Internet-of-Things. We have in this paper investigated trust in an IoT setting in considerable depth, and while it is obvious that one cannot fully trust any of the IoT components (software, hardware, communications, etc), it does not mean that humans cannot or should not trust IoT services at all. The human heuristic handling of risks, threats and opportunities is not without its faults, but use of trusted proxy devices and the trust we have in recognized brands and companies will enable us to trust many services without too much hesitation.

From a theoretical perspective, we have seen that there are complexities and that the dynamic nature of IoT environments are hard to capture and difficult to model accurately. An inherent part of the problem is that humans have no way of ascertaining the true intent of the device. Furthermore, humans lack methods for determining the ability of the device to behave as expected. However, we are optimistic that approaches like TNA-SL will prove useful for modeling human-to-device trust in an IoT settings.

References

1. Thompson, K. (1984). Reflections on trusting trust. Turing award lecture. *Communications of the ACM*, 27(8), 761–763.
2. Raymond, E. S. (1999). *The Cathedral and the Bazaar*. USA: O'Reilly Media.
3. Anderson, R. (2005). Open and closed systems are equivalent (that is, in an ideal world). In J. Feller, B. Fitzgerald, S. A. Hissam, & K. R. Lakhani (Eds.), *Perspectives on free and open source software*. Cambridge: MIT Press.
4. ISO, ISO/IEC 11990-1. (2009). Information technology: Trusted platform module—Part 1: Overview.
5. 3GPP. (2011). TS 31.101 UICC-terminal interface; Physical and logical characteristics (Release 9), 01-2011.
6. 3GPP. (2010). TS 33.320 security of home node B (HNB) / home evolved node B (HeNB) (Release 11), 12-2010.
7. Chakraborty, R. S., Narasimhan, S. & Bhunia, S. (2009). Hardware Trojan: Threats and emerging solutions. In *Proceedings of IEEE International High Level Design Validation and Test Workshop 2009 (HLDVT'09)* (pp. 166–171).
8. Tehranipoor, M. & Koushanfar, F. (2010). A survey of hardware Trojan taxonomy and detection. In *Design & Test of Computers*, 27(1), 10–25, IEEE Computer Society.
9. Jin, Y., Kupp, N. & Makris, Y. (2009). Experiences in hardware Trojan design and implementation. In *Proceedings of IEEE International Workshop on Hardware Oriented Security and Trust, 2009, (HOST'09)* (pp. 50–57). August 2009.
10. Falliere, N., Murchu, L. O. & Chien, E. (2011). W32.Stuxnet Dossier. Symantec, Version 1.4, February 2011.
11. Køien, G. M. (2009). *Entity authentication and personal privacy in future cellular systems*. Denmark: River Publisher.
12. Liu, Y. & Wang, K. (2010). Trust control in heterogeneous networks for internet of things. In *Proceedings of IEEE International Conference on Computer Applications and Systems Modelling (ICCASM) 2010* (pp. 632–636). IEEE Press, October 2010.
13. Grandison, T. & Sloman, M. (2000). A survey of trust in internet applications. *Communications Surveys & Tutorials*, 3(4), 2–16. IEEE Communications Society.
14. Dolev, D. & Yao, A. (1983). On the security of public key protocols. *IEEE Transactions on Information Theory*, 29(2), 198–208. IEEE Information Theory Society, March, 1983.
15. West, R. (2008). The psychology of security: Why do good users make bad decisions? *Communications of the ACM*, 51(4), 34–40.
16. Schneier, B. (2003). *Beyond fear: Thinking sensibly about security in an uncertain world*. New York: Springer.
17. Baron-Cohen, S. (2003). *The essential difference: Men, women and the extreme male brain*. London: Penguin.
18. Dawkins, R. (2006). *The selfish gene*. Oxford University Press, England. 30th anniversary edition.

19. Dawkins, R. (2008). *The extended phenotype*. Oxford University Press, England. Revised and corrected edition.
20. Robinson, J., Wakeman, I., Chalmers, D. & Horsfall, B. (2010). Trust and the internet of things. In *Proceedings of TruLoco 2010: The Joint International Workshop on Trust in Location and Communications in Decentralised Computing*.
21. Baars, B. J., & Gage, N. M. (2010). *Cognition, brain, and consciousness: Introduction to cognitive neuroscience* (2nd ed.). London: Academic Press.
22. Goleman, D. (2006). *Emotional intelligence: Why it can matter more than IQ*. 10th Anniversary Edition, New York: Bantam. September 2006.
23. Jøsang, A., Hayward, R. & Pope, S. (2006). Trust network analysis with subjective logic. In *Proceedings of the 29th Australasian Computer Science Conference (ACSC'06)*, Australian Computer Society.
24. Jøsang, A. & Bhuiyan, T. (2008). Optimal trust network analysis with subjective logic. In *Proceedings of the Second International Conference on Emerging Security Information, Systems and Technologies (SECUREWARE'08)*. IEEE Computer Society, 2008.

Author Biography



Geir M. Kjøien holds an associate professor position at the University of Agder, Norway, where he works with cellular/wireless systems, access security, security protocols, formal verification, privacy, trust and trustworthiness, and similar topics. He holds a B.Sc. Hons in Computing Science from the University of Newcastle upon Tyne, England, an M.Sc. in IT from the Norwegian University of Science and Technology (NTNU, then NTH), and a Ph.D. from Aalborg University (AAU). He is a senior member of ACM and a senior member of IEEE. Kjøien has previously worked for LM Ericsson Norway, System Sikkerhet AS and Telenor R&D, and participated in security standardization in 3GPP during 1999–2009 as the Telenor delegate. He has worked extensively with access security in GSM/GPRS, UMTS and LTE.