Tailored Security and Safety for Pervasive Computing

Erik-Oliver Blass[1] and Martina Zitterbart[2]

[1]EURECOM, Sophia Antipolis, France
[2]Institut für Telematik, Universität Karlsruhe (TH), Germany

Abstract. Pervasive computing makes high demands on security: devices are seriously resource-restricted, communication takes place spontaneously, and adversaries might control some of the devices. We claim that 1.) today's research, studying traditional security properties for pervasive computing, leads to inefficient, expensive, and unnecessary strong and unwanted security solutions. Instead, security solutions tailored to the demands of a user, the scenario, or the expected adversary are more promising. 2.) Today's research for security in pervasive computing makes naive, inefficient, and unrealistic assumptions regarding safety properties, in particular the quality of basic communication. Therefore, future security research has to consider safety characteristics and has to jointly investigate security and safety for efficient, tailored solutions.

1 Introduction

Soon, tiny computer systems will surround us in large numbers and help us mastering our everyday life. For example, small sensors integrated into patients' clothes will continuously monitor the state of health, intelligent RFID tags will allow for tracking of objects or identifying groups of persons, and in our homes, sun blinds, air condition and lighting will autonomously coordinate themselves. In general, physically small devices will autonomously form networks and communicate wirelessly, unnoticed by the user. In such a vision of pervasive computing, many new security, privacy, and safety challenges arise.

2 New Challenges

Providing security for pervasive computing is difficult: as small devices typically feature only simplistic CPUs due to size and cost restrictions, classical security solutions become unfeasible. Communication between devices takes place spontaneously, central infrastructure components, such as key-distribution servers or Certificate Authorities, are not available, impeding establishment of trust. Of capital importance, however, is energy: tiny devices are battery-powered. So, after their energy is depleted, they cannot offer their service anymore. Energy-expensive cryptographic computations, frequent wireless communication, as with

current security solutions, will deplete batteries and quickly render devices useless. Today's research in security tackles these problems and already investigates solutions.

This paper identifies the following open research problems:

1. **Tailor Security.** Today's research focuses on implementing traditional, strong security guarantees, all-or-nothing security guarantees, for pervasive computing. We claim that such strong security guarantees will not be affordable in many pervasive computing scenarios. Instead, better-than-nothing, tailored security solutions, tailored to the demands of the user and the capabilities of the system, are more suitable. They will lead to compromise between security guarantees and, e.g., energy consumption.
Also, especially in pervasive computing, the capabilities of the *adversary* have to be considered. Contrary to traditional adversary models, e.g., Dolev and Yao [5], stronger adversaries have to be assumed that might not only eavesdrop communication, but also easily take over some of the tiny, typically unprotected devices or even add malicious devices to a network [1]. To protect against stronger adversaries, additional security means have to be provided. Coping with stronger adversary requires more expensive protocols. Finally, new security properties such as un-traceability and privacy have to be taken into account.
2. **Integrate Safety.** Today's research in security often makes naive assumptions with respect to safety, robustness, and reliability. For example, security research often assumes error-free communication as given by lower layers. We claim that security just on top of safety is inefficient and counterproductive. Future research will have to integrate security and safety and aim at a combined framework for pervasive computing.

3 Tailored Security

In future research, one should analyze, design, and implement non-classical, better-than-nothing security solutions for pervasive computing. As traditional solutions for typical security properties such as confidentiality, integrity, and availability cannot be applied, we propose a new security paradigm that we call *tailored security.* To cope with the special properties of pervasive computing scenarios and especially to minimize energy-consumption for extension of devices' lifetimes, we propose to trade-off security properties against energy consumption. In an environment where classical, but energy-expensive strong security properties against a strong adversary cannot be afforded or are not required, the user of a service might accept, or might only be in the position to afford relaxed, i.e., weaker security properties. If users accept weaker security, they can benefit from an extended lifetime of their devices.

User adjustable. New adjustable security solutions should be designed, which can be parameterized by the user according to his needs in terms of security and

available budget in terms of, e.g., remaining battery-power. Generally, if the user is willing to accept weaker security properties due to his needs and requirements, he will get something beneficial back in return, for example, aforementioned higher network lifetime. One could imagine the user to opt for a high level of security, if a lot of energy is available, or opt for weaker security to achieve maximum network lifetime.

Self-Governed. Additionally, the "system" might, with respect to some user's guidelines, automatically adapt security over time. So, security levels might gradually decrease as the batteries of devices start to deplete. If batteries are refreshed or recharged, security levels and therewith energy consumption might rise again. While such behavior might be exploited by an adversary, i.e., by waiting until security levels are low enough to break, this still provides security for as long as possible. To cope with different adversaries, future solutions should monitor adversarial behavior and, e.g., using feedback mechanisms or self-monitoring, adjust security levels, if malicious behavior is detected in the system. Also, in case of a DoS-attack, the system could gradually decrease services to protect against battery depletion, and to provide as much security for as long as possible.

Finally, there will be security properties where gradual weakening is impossible: for example, without physically compromising a device, an adversary should not be able to recover the device's master secret.
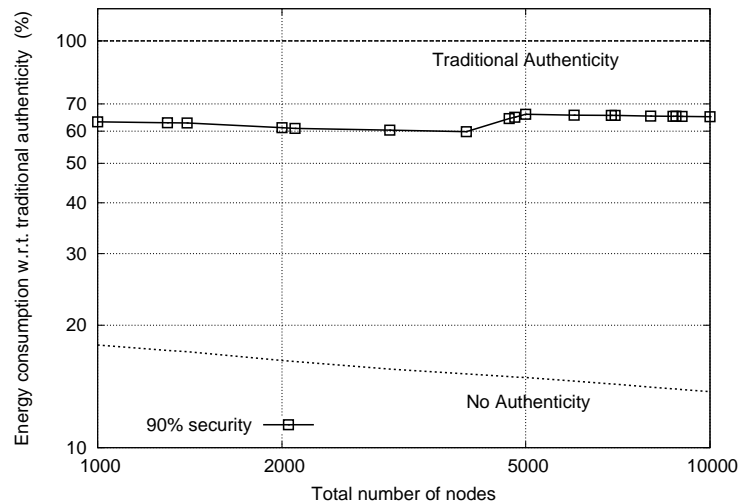
3.1 First Steps

Fig. 1. Tailoring authenticity to save energy against 20% compromised devices [3]

How to degrade security. First steps towards weakening security levels, a possibility to extend devices' lifetime, could be a probabilistic relaxation of security against a certain adversary. Probabilistic relaxation of authenticity to save energy for data transport in wireless sensor networks has been proposed in preliminary work [3]. With probabilistic relaxation, the main idea would be to accept a probabilistic level of security: for example, $\approx 90\%$ of all measurements in a sensor network are secure against an adversary being able to compromise 20% of all devices. This relaxations should save a large amount of energy compared to a traditional "100%" security solution. Figure 1 illustrates this idea by showing simulation results of using a data transport protocol within a wireless sensor network and guaranteeing only probabilistic authenticity against a certain fraction of compromised devices (i.e., devices under control of the adversary). In particular, energy savings for authentic data transport are depicted. The x-axis shows the number of devices in the sensor network, while the y-axis shows the per device energy consumption. Energy consumption is, however, not given in total amounts, but as a percentage of the energy consumption that would be necessary for traditional authenticity. For comparison, the lower bound for energy consumption, i.e., no authenticity at all, is shown in Figure 1. You can see from Figure 1 that a relaxed authenticity of "90%" already allows for more than 30% of energy savings compared to traditional security in the presence of 20% compromised devices. On the one hand, lower authenticity requirements will result in higher energy savings. On the other hand, higher authenticity requirements or protection against a larger fraction of compromised devices will increase energy consumption [3].

Applying probabilistic relaxation of security, the user can tailor his system's security to his needs. Future research should extend and generalize the idea of probabilistic relaxation to all possible security properties, e.g., confidentiality.


Fig. 2. Tailoring security properties [12]

In addition, another approach would be to graduate whole security properties, see Figure 2. A system that simply detects malicious behavior in a network, offering *Completeness*, will be "cheaper" than a system that does *both*, detects misbehavior and "compensates" for it, thus offering *Soundness*. A system that can detect *and* compensate malicious behavior in the network, *and* that can even identify the responsible malicious devices, as with *non-repudiation*, will be even more expensive. The user chooses security properties he wants and can afford.

Note that previous work, cf., Lindskog and Brunstrom [9], Lindskog et al. [10], focuses on "tuning" IPSec parameters and selectively encrypting packets to achieve different performance results or energy costs. Here, the user can chose between AH or ESP modes and chose the cryptographic primitives, such as MD5 or SHA-1 and DES or 3DES. Also, the length of cryptographic keys and the number of encryption rounds can be reduced, see Chandramouli et al. [4], Irvine and Levin [7]. While these are clearly aspects of tailored security as proposed in this paper, it is difficult to analyze the security difference between different cryptographic primitives or varying the number of encryption rounds. More importantly, there is no notion of different adversarial capabilities and the threat of compromised devices, unique in pervasive computing scenarios. Also note that, e.g., *Soundness* in Figure 2 comprises *Completeness*, and *Non-Repudiation* comprises *Soundness*. We, therefore, do not tailor "orthogonal" security properties such as confidentiality and authenticity as Irvine and Levin [6]: it is difficult to find a metric or taxonomy within orthogonal security properties.

How to define a Trade-off. Given the mechanisms to weaken or trade-off security, another important research issue is to investigate how users can specify trade-offs between security and something beneficial in return. Two possible first steps in defining or tailoring trade-offs to the user's demands can already be identified as follows: either explicitly by the user, by setting security or energy parameters with respect to, e.g., the demanded lifetime of his system or the energy he is willing to spend for security. Using theoretical or practical estimates, the user can explicitly tailor his security. The second way would be the aforementioned automatic way of setting security levels, where the system itself tries to deliver the best possible security for as long as possible following some guidelines, feedback, or rules given by the user in advance.

Consider new Security Properties. Besides traditional security properties, future research must integrate many new, more recent security aspects. If the user is surrounded by tiny devices, especially privacy of his data becomes more and more important. So, data should not only be, e.g., confidential, but the adversary must also not be in the position to deduce the owner of the data, data origin and destination, or the type of data. Besides being anonymous in the system, the user actions or user data should also not be linked or traced. Even if the owner of data cannot be identified by an adversary he should also not be able to link different transactions or data seen to the same origin. While there is already some research going on in this area, solutions again focus on providing perfect privacy, i.e., perfect "1-out-of-n" indistinguishability and unlinkability of all devices under all circumstances, cf., Juels and Weis [8]. Once more, this kind of privacy might either not be affordable due to the requirement of complex cryptographic operations or is not required due to the demands of the user. In many scenarios, it might be even sufficient to offer even only "1-out-of-2" indistinguishability, so cheaper trade-of based solutions can again be investigated. In

conclusion, new security aspects, such as privacy, should also be tailored to the user's demands and integrated into today's security protocols.
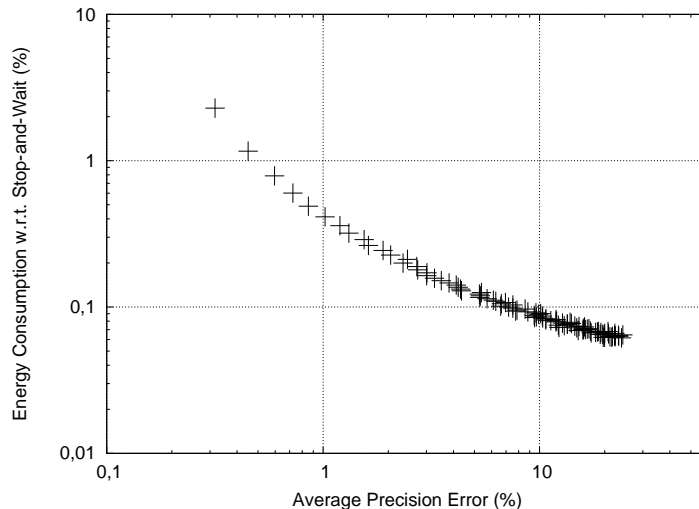
4   Integrate Safety

Fig. 3. Prevision vs. energy trade-off

Additionally, we propose to consider and integrate aspects of safety or robustness into the security context. For example, today's research in security typically requires and thus assumes perfect, error-free communication between devices. This is totally unrealistic in pervasive computing environments, as wireless communication using cheap radio interfaces is prone to errors, and end-to-end connectivity cannot be guaranteed. Usage of classical ARQ techniques, such as Stop-and-Wait, as a basis to guarantee robust communication and then building security on top of it, is often assumed, however very inefficient in terms of energy consumption. For example, with cheap and simplistic radio interfaces 50% and more packet-loss can easily occur, cf., Turau et al. [11]. If data $x_1, x_2, \ldots$ needs to be send from a sender to a receiver, using Stop-and-Wait with 50% pack-loss leads on average to a total of 6 costly transmissions per date x_i – which is way to energy-inefficient.

First, as with security properties, some of the traditional safety properties, in-order delivery, duplicate freeness, and error-freeness, might become superfluous. In a scenario where a body area network monitors the average body temperature of a patient, in-order delivery is nonessential. Yet, new application specific parameters such as the precision or the freshness of data processed in the system become important.

In preliminary work, we have shown that reducing the precision of data transported between devices by a small percentage allows huge energy savings [2]. Figure 3 depicts the energy consumption for a protocol sending sensor measurements from a sender to a receiver. This protocol, however, does not send these measurements error-free. Instead, a certain error in the precision of each measurement sent is allowed. The x-axis show the average precision error for each measurement, the y-axis show the resulting energy consumption. Again, this energy consumption is not given as a total amount, but as a percentage of the energy consumption that would be necessary to send each measurement error-free (using Stop-and-Wait, 50% packet loss). The precision error is given as $\frac{|x_{\text{sent}} - x_{\text{received}}|}{|x_{\text{max}} - x_{\text{min}}|}$, where x_{sent} is the value of the measurement as sent, x_{received} is the value of the measurement as received, x_{max} and x_{min} represent the maximum and minimum values measurements can take during a certain time period.

Again, the idea of percentage weakening can be generalized and extended to all aspects of safety.

Second, future research should jointly investigate and combine security and appropriate safety mechanisms. Security solutions should be able to cope with device or transmission failures. For example, failure of transmissions of encrypted data could be simply accepted by the system. Similar to approaches of today's video and audio codecs, any "missing" or "faulty" data could be interpolated from data received properly before. The challenge will be to find solutions suitable for resource-restricted hardware of typical pervasive computing devices.

In case of high-packet loss, data could be sent redundantly, but without HMACs for every single packet integrity could now be verified by one single, combined HMAC sent with ARQ techniques. This saves both, costly computations of per-packet HMACs as well as per-packet ARQ communication.

5 Conclusion

Research in pervasive computing security has so far mostly focused on implementing traditional security properties into new environments. Furthermore, all safety or robustness aspects are considered as building blocks, and security solutions are building on top of it. Instead, we propose future research to investigate two new areas: 1.) Research should jointly analyze security and safety to develop energy efficient solutions, and 2.) future research should aim at designing tailored solutions, suited to the user's needs and the devices' capabilities.

Bibliography

[1] Z. Benenson, P.M. Cholewinski, and F.C. Freiling. *Wireless Sensor Network Security*, volume 1, chapter Vulnerabilities and Attacks in Wireless Sensor Networks, pages 22–43. IOS Press, 2008. ISBN 978-1-58603-813-7.

[2] E.-O. Blass, L. Tiede, and M. Zitterbart. An energy-efficient and reliable mechanism for data transport in wireless sensor networks. In *Proceedings of Third International Conference on Networked Sensing Systems*, pages 211–216, Chiacgo, USA, 2006. ISBN 0974361135.

[3] E.-O. Blass, J. Wilke, and M. Zitterbart. Relaxed authenticity for data aggregation in wireless sensor networks. In *Proceedings of Fourth International Conference on Security and Privacy in Communication Networks*, pages 1–10, Istanbul, Turkey, 2008. ISBN 9781605582412.

[4] R. Chandramouli, S. Bapatla, K.P. Subbalakshmi, and R.N. Uma. Battery power-aware encryption. *ACM Transactions on Information and System Security*, 9(2):162–180, 2006. ISSN 1094-9224.

[5] D. Dolev and A. Yao. On the security of public key protocols. *IEEE Transactions on Information Theory*, 29(2):198–208, 1983. ISSN: 0018-9448.

[6] C. Irvine and T. Levin. Toward a taxonomy and costing method for security services. In *Proceedings of Annual Computer Security Applications Conference*, pages 183–188, 1999. ISBN 0-7695-0346-2.

[7] C. Irvine and T. Levin. Quality of security service. In *Proceedings of New Security Paradigms Workshop*, pages 91–99, Ballycotton, Ireland, 2001. ISBN 1-58113-260-3.

[8] A. Juels and S. Weis. Defining strong privacy for rfid. In *Proceedings of PerCom Workshops*, pages 342–347, White Plains, USA, 2007. ISBN 9-780-769-527-888.

[9] S. Lindskog and A. Brunstrom. Design and implementation of a tunable encryption service for networked applications. In *Proceedings of 1st workshop on Security and Privacy for Emerging Areas in Communication Networks*, pages 258–266, 2005. ISBN 0-7803-9468-2.

[10] S. Lindskog, Z. Faigl, and A. Brunstrom. A conceptual model for analysis and design of tunable security services. *Journal of Networks*, 3(5):1–12, 2008.

[11] V. Turau, C. Renner, M. Venzke, S. Waschik, C. Weyer, and M. Witt. The heathland experiment: Results and experiences. In *REALWSN: Real World Wireless Sensor Networks*, 2005.

[12] J. Wilke, E.-O. Blass, F.C. Freiling, and M. Zitterbart. A framework for probabilistic, authentic aggregation in wireless sensor networks. *PIK – Praxis der Informationsverarbeitung und Kommunikation*, 32(2), 2009. ISSN 0930-5157 (to appear).