

HELSINKI UNIVERSITY OF TECHNOLOGY  
Department of Electrical and Communications Engineering

# User centred access control for home networks

Master's Thesis

RELEASE CANDIDATE 1-p.1, v RC1-p1-0701 @ July 1, 2007

Robin Laurén

Telecommunications Software and Multimedia Laboratory  
Espoo 2007



|                       |  |
|-----------------------|--|
| <b>Author:</b>        | Robin Laurén   |
| <b>Title:</b>         | User centred access control for home networks  |
| <b>Date:</b>          | July 1 <sup>st</sup> , 2007  |
| <b>Language:</b>      | English  |
| <b>Pages:</b>         | 80 (+ approx. 20)  |
| <b>Professorship:</b> | T-110 Telecommunications software and multimedia laboratory (TML)  |
| <b>Supervisor:</b>    | Professor Antti Ylä-Jääski   |
| <b>Instructor:</b>    | Ursula Holmström, M.Sc.  |
| <b>Abstract:</b>      | <p>It is anticipated that the next development in domestic technology is to have the devices interoperate in a networked fashion, much like office technology currently does. We call this interoperating, ubicomp-like domestic technology the Home Network. The security of home networks, and specifically <i>access control for home networks</i>, is the focus of this thesis – an area that yet hasn't received much attention neither by academic or commercial camps.</p> <p>In this work, we review what a home network is, what means exist for access control that are suitable for a home deployment, and which ways not specifically intended for domestic use would be adaptable to the home network. Five families were interviewed for this work, each with children, to give us some indication on how families view and use technology, and use their spare time. We based our interviews on the Contextual design method. Also, security experts were interviewed to share their insight on what should be relevant for home network security and access control.</p> <p>A central finding is that access control for users within the home can largely be handled with social norms or 'the social barrier'. To reflect this behaviour, a Door metaphor is presented where a user is allowed to make a decision whether to breach the privacy requested by its owner, with the suggested addition that all accesses through the Door are logged and log entries are viewed at the owner's discretion. The behaviour observed in the families regarding access control was coherent with the Generalized role based access model (GRBAC). Therefore it is suggested that a system to help an administrative user to build a security policy for the home network would be based on the GRBAC model.</p> |
| <b>Keywords:</b>      | Usability, security, HCIsec, user-centred security, home networks, domestic technology, ubiquitous computing, pervasive computing, access control.   |

The thesis is available on-line at <http://robin.lauren.fi/thesis>. The author can be contacted at Robin@Lauren.fi.

|                        |  |
|------------------------|--|
| <b>Författare:</b>     | Robin Laurén   |
| <b>Arbetets titel:</b> | User centred access control for home networks (Användarcentrerad tillgänglighetskontroll för hemmanätverk)   |
| <b>Datum:</b>          | Den 1 juli 2007  |
| <b>Språk:</b>          | Engelska   |
| <b>Sidantal:</b>       | 80 (+ ca. 20)  |
| <b>Professur:</b>      | T-110 Laboratoriet för telekommunikationsprogram och multimedia (TML)  |
| <b>Övervakare:</b>     | Professor Antti Ylä-Jääski   |
| <b>Handledare:</b>     | DI Ursula Holmström  |
| <b>Sammandrag:</b>     | <p>Hemelektronik och övrig teknisk apparatur förväntas inom nära framtid utvecklas så att hemtekniken kan börja samarbeta i ett nätverk jämförbart med nätverk på arbetsplatser. Dessa hemteknikens system kallas i detta arbete för hemnätverk. Säkerhetsaspekter specifikt för hemnät, och speciellt <i>tillgänglighetskontrollen för hemmanätverk</i>, är ämnen som ännu inte givits mycket uppmärksamhet i vare sig den akademiska eller den kommersiella världen.</p> <p>I detta diplomarbete beskrivs vad som kan avses med hemnätverk, vilka metoder för tillgänglighetskontroll som kan anses passliga i en hemnätsmiljö, och vilka metoder som kan anpassas för hemnätbruk. För arbetet intervjuades fem banrfamiljer om hur de uppfattade hemtekniken, säkerhet och tillgänglighetskontroll, hur de använde datateknik på fritiden som stöd för deras fritidssysslor. Intervjuerna och deras analys gjordes baserat på Contextual design-metoden. Experter inom datatekniken intervjuades också, för att ge vidare insikt om vad de ansåg viktigt med tanke på säkerhet och tillgänglighetskontroll i hemmanät.</p> <p>En upptäckt var att de sociala mönstren och gränserna på vad som är tillåtet och otillåtet sträckte sig även till hemtekniken. Sålunda föreslås att tillgänglighetskontrollen dels kan basera sig på dessa mönster i form av en Dörr-metafor. Enligt Dörrmetaforen tillåts de övriga invånarna göra ett etikbaserat beslut över huruvida de anser det vara godkännbart eller ej att ta sig igenom Dörren.</p> <p>En annan upptäckt var att familjerna följde den Allmänna rollbaserade tillgänglighetskontrollmetoden (GRBAC). Sålunda rekommenderas att ett system som ska hjälpa personen eller personerna bygga hemnätets säkerhet och tillgänglighetskontroll baseras just på GRBAC-modellen.</p> |
| <b>Nyckelord:</b>      | Användbarhet, säkerhet, HCIsec, användarcentrerad säkerhet, tillgänglighetskontroll, hemnätverk, hemteknik, ubiquitous computing (överallt förekommande datateknik).   |

Detta diplomarbete finns tillgängligt på adressen <http://robin.lauren.fi/thesis>. Författaren kan nås per epost på adressen Robin@Lauren.fi.

|                     |   |
|---------------------|---|
| <b>Tekijä:</b>      | Robin Laurén  |
| <b>Työn nimi:</b>   | User centred access control for home networks (Kotiverkkojen käyttäjäkeskeinen pääsynvalvonta)  |
| <b>Päiväys:</b>     | July 1 <sup>st</sup> , 2007   |
| <b>Kieli:</b>       | English   |
| <b>Sivumäärä:</b>   | 80 (+ n. 20)  |
| <b>Professuuri:</b> | T-110 Tietoliikenneohjelmistjen ja multimedian laboratorio (TML)  |
| <b>Valvoja:</b>     | Professor Antti Ylä-Jääski  |
| <b>Ohjaaja:</b>     | DI Ursula Holmström   |
| <b>Yhteenveto:</b>  | <p>Kotikäyttöön tarkoitettu teknologia oletetaan lähitulevaisuudessa kehittyvän niin että kodin eri laitteet saadaan verkottumaan keskenään, samaan tapaan kuin toimistojen tietotekniikka nyt. Tällaista kodin jokapaikan tietotekniikkaa käsitellään tässä työssä nimellä Kotiverkko. Nimenomaan kotiverkkojen tietoturva, ja erityisesti <i>kotiverkkojen pääsynvalvontaa</i> ei ole liioin käsitelty tieteellisessä, saati kaupallisessa kirjallisuudessa.</p> <p>Tässä diplomityössä esitellään mitä kotiverkko on, mitä pääsynvalvonnan menetelmiä löytyy jotka ovat omiaan juuri kotiverkoissa, sekä mitä muissa yhteyksissä esiteltyjä menetelmiä voisi muuntaa käytettäväksi kotiverkoissa. Keskeinen löydös oli että kotiverkoissa pääsyä rajoittaa eteenkin hyvät tavat ja sosiaaliset normit. Sen perusteella esitellään käytettäväksi metafora Ovesta, josta muut kuin oven omistaja voivat päättää jos sen läpikulkeminen on omistajan toiveiden mukaan hyväksyttävää. Kaikki kulut voidaan kirjata lokiin ja tarvittaessa tai pyydettyäessä esittää omistajalle.</p> <p>Osana työtä haastateltiin viittä lapsiperhettä, tarkoituksena antaa kuva siitä miten ihmiset kokivat tietotekniikan ja käyttivät sitä vapaa-ajallaan sekä vapaa-ajantoiminnan tukena. Haastattelut sekä niiden analysointi tehtiin Contextual design-menetelmään perustuen. Myös tietotekniikan asiantuntijoita haastateltiin saadakseen heidän näkemyksensä kotiverkkojen tietoturva- ja pääsynvalvontatarpeista.</p> <p>Kodeissa yllämainitun sosiaalisiin normeihin perustuvan kulunvalvonnan lisäksi perheiden käyttäytymisissä oli selviä piirteitä Yleistetystä rooliperusteisesta pääsynvalvontamenetelmästä (GRBAC). Tähän huomioon perustuen esitetään että työkalu kotiverkon pääsynvalvonnan kuvaamiseksi perustuisi juuri GRBAC-malliin.</p> |
| <b>Avainsanat:</b>  | Käytettävyys, tietoturva, käyttäjälähtöinen tietoturva (HCIssec), kotiverkot, pääsynvalvonta, kodin tietotekniikka, jokapaikan tietotekniikka, saavutettavuus.  |

Tämä diplomityö löytyy Internetistä osoitteesta <http://robin.lauren.fi/thesis>. Tekijän sähköpostiosoite on Robin@Lauren.fi.

## Suggested BibTeX reference

```
@mastersthesis{ Lauren2007,  
  author = "Robin Laur'en",  
  title = "User centred access control for home networks",  
  school = "Helsinki University of Technology (TKK)",  
  year = "2007",  
  url = "http://robin.lauren.fi/thesis"  
}
```

Created through extensive use of Kile, Emacs, KDissert, Dia, Inkscape, The Gimp, Google Scholar, The Dividing Line, and a host of other programs, services (yes, including Wikipedia), potables and consumables. Typeset with L<sup>A</sup>T<sub>E</sub>X. Source code available upon request. (CC) BY-NC-SA.

# Contents

|  |            |
|--|------------|
| <b>Abstract</b>  | <b>iii</b> |
| Sammandrag . . . . .                                       | iv         |
| Yhteenveto . . . . .                                       | v          |
| <b>Acknowledgements</b>                                    | <b>xii</b> |
| <b>1 Introduction</b>                                      | <b>1</b>   |
| <b>2 Problem statement and criteria</b>                    | <b>5</b>   |
| 2.1 Problem statement . . . . .                            | 6          |
| 2.2 Scope . . . . .  | 7          |
| <b>3 The home network</b>                                  | <b>9</b>   |
| 3.1 What is a home network? . . . . .                      | 9          |
| 3.2 Terminology soup . . . . .                             | 16         |
| 3.3 Requirements for a home network . . . . .              | 17         |
| 3.4 Homes and their ‘users’ . . . . .                      | 18         |
| 3.5 Stakeholders . . . . .                                 | 19         |
| 3.6 Backline technology: Universal Plug and Play . . . . . | 21         |
| <b>4 Usable security</b>                                   | <b>27</b>  |
| 4.1 Usability and User centred design . . . . .            | 28         |
| 4.2 Contextual Design . . . . .                            | 30         |
| 4.3 Security basics . . . . .                              | 33         |
| 4.4 Access control . . . . .                               | 35         |
| 4.4.1 Motivation . . . . .                                 | 38         |

|          |  |           |
|----------|--|-----------|
| 4.4.2    | Means and methods . . . . .                              | 41        |
| 4.4.3    | Multi-factor authentication . . . . .                    | 43        |
| 4.5      | Making security usable . . . . .                         | 43        |
| <b>5</b> | <b>Previous work</b>                                     | <b>47</b> |
| 5.1      | Security on home computer networks . . . . .             | 47        |
| 5.2      | Identified problems on pervasive networks . . . . .      | 48        |
| 5.3      | User identification and authentication . . . . .         | 50        |
| 5.3.1    | Context based authentication . . . . .                   | 51        |
| 5.4      | Device management . . . . .                              | 52        |
| 5.5      | Role-Based Access Control and GRBAC . . . . .            | 54        |
| 5.6      | Security mark-up . . . . .                               | 55        |
| <b>6</b> | <b>Access control at home</b>                            | <b>57</b> |
| 6.1      | User interviews . . . . .                                | 57        |
| 6.1.1    | Findings . . . . .                                       | 60        |
| 6.2      | Expert interviews . . . . .                              | 63        |
| <b>7</b> | <b>Analysis</b>  | <b>69</b> |
| 7.1      | Family interview findings . . . . .                      | 69        |
| 7.2      | Comparing family and expert opinions . . . . .           | 71        |
| 7.3      | A sketch for a usable access control ecosystem . . . . . | 71        |
| 7.3.1    | The Door: responsibility driven access control . . . . . | 74        |
| 7.3.2    | Other security ideas . . . . .                           | 76        |
| 7.4      | The applicability of Contextual design . . . . .         | 77        |
| 7.5      | For future research . . . . .                            | 77        |
| <b>8</b> | <b>Conclusion</b>  | <b>79</b> |

# List of Figures

|     |  |    |
|-----|--|----|
| 3.1 | The home network environment . . . . .                           | 10 |
| 3.2 | A somewhat fictional home network plan . . . . .                 | 14 |
| 3.3 | Abraham Maslow's Hierarchy of needs . . . . .                    | 17 |
| 4.1 | The facets of acceptability . . . . .                            | 28 |
| 4.2 | Floor plan of typical (but hypothetical) usability lab . . . . . | 29 |
| 6.1 | Our affinity wall in mid-progress . . . . .                      | 59 |
| 7.1 | A hypothesis: User control versus user happiness . . . . .       | 72 |

## Acronyms and abbreviations

Acronyms and their usage in the context of this thesis. Some of these acronyms have other meanings outside the scope of this text.

|       |   |
|-------|---|
| ACL   | Access control list, lists wich actors (or groups of actors) have what access rights to a given resource.                       |
| CD    | Contextual Design.  |
| DAC   | Discretionary Access Control, a type of access control which is administered by the owners of each resource (see MAC, sense 1). |
| DCP   | Device Control Protocol, a description of the protocol of a UPnP device.  |
| DRM   | Digital rights management, a way to control the use and distribution of digital media.  |
| GRBAC | Generalized Role-Based Access Control. See RBAC below, and § 5.5.   |
| HAN   | Home Area Network, a domestic network of computers and digital devices.   |

|        |   |
|--------|---|
| HCIsec | The married field of usability (Human Computer Interfacing) and security.   |
| ISP    | Internet service provider.  |
| MAC    | <b>1.</b> Mandatory Access Control, a type of access control which is centrally administered. Not to be confused with <i>Media Access Control (MAC) Address</i> , the unique identifier of (each interface of) a networked device – despite its name not an access control mechanism at all. <b>2.</b> Message Authentication Code, a “secret value” known by parties engaging in symmetrically encrypted communication.  |
| PIN    | Personal identification number (often called <i>PIN code</i> ), a usually four digit password.  |
| RFID   | Radio frequency identification, a technology for transmitting and receiving user identification data over a short range wireless link.  |
| RBAC   | Role-Based Access Control, a type of access control based on the role (or roles) of a user.   |
| UCD    | User centred design, a design philosophy and a process centred on needs, wants, and limitations of the end user.  |
| UI     | User Interface, the “visible” parts of an appliance or application.   |
| UPnP   | Universal Plug and Play, a family of protocols to allow devices to connect and interoperate.  |
| UTP    | Unshielded Twisted Pair, a common type of network cabling where two conductors carrying equal and opposite signals are wound together to cancel out electromagnetic interference from external sources and neighbouring wires (cf. FTP, Foiled Twisted Pair which has a foil shielding around the wires, and STP, Shielded Twisted Pair which has each wire pair shielded; both techniques add protection to/from external signals). Most computer network cables in use in 2007 are of UTP type. |
| UUID   | Universally Unique Identifier, also known as GUIDs (Globally unique identifiers, especially in Microsoft lingo). A 128-bit “fingerprint” which can be guaranteed to be unique in space and time [IT03, LMS05]. Canonically expressed in a hexadecimal form such as 550e8400-e29b-41d4-a716-446655440000.  |

## Glossary

|               |   |
|---------------|---|
| backline      | A more relaxed term for infrastructure.   |
| domotics      | Domestic informatics, the application of information technology and robotics in a domestic setting.   |
| node          | Any addressable component on a network. Devices, appliances, sensors and computers are all nodes. Also known as network <i>endpoints</i> in e.g. UPnP.                |
| payload       | The contents of a transmission when metadata such as headers and checksums has been stripped away.  |
| peanut device | A small, wireless, battery-driven nodes on the network, low on computational power.   |
| principal     | Any entity that can perform actions on the network; encompassing, without distinction, humans, machines acting as representatives of humans, and machines that don't. |
| ubiquitous    | Present or appearing everywhere; omnipresent.   |

## Acknowledgements

Thanks to Antti Ylä-Jääski for having me at his lab, Ursula Holmström for helping me with the thesis, and all members of the TML lab for input and inspiration. Thanks and apologies to all who stood by and endured the maturing process of this work and my long-winding path to examination.

Special thanks to Hanna for shoving me up that hill. And for Ronja and Linus. You are the wind at my back [Mor02].

Kudos to Carl Ellison and Frank Stajano for their inspiring work,<sup>1</sup> to Rene Mayrhofer for help and feedback, and to Janelle Saffin and the Australian Government for borrowing me a laptop in Timor-Leste, where the work on this thesis begun.

I dedicate this work to my father, Christian Laurén, who taught me that words are fun and who i still dearly miss after 25 years.

In Helsinki, July 1, 2007

– Robin Laurén

---

<sup>1</sup>I am buggered that i didn't read Stajano's book "Security for Ubiquitous Computing" early on in the process, as it would have helped me finding the direction for this thesis immediately. It's a great book if you're at all into the kind of material covered here. And without Ellison's "Home network security", this thesis would never even have been thought of.

# Chapter 1

## Introduction

Families use computers increasingly for a variety of recreational activities [GEND05]. Computers, once scientific tools and office machines, are used for shopping, information gathering, learning and communications, gaming and music playing [VSSM01]. Complementing the computers are portable digital music players and digital cameras. Even household appliances such as washing machines are equipped with microprocessor technologies, and although few home users think of them as specialized computers, they are.

There are two major, and interlinked, developments on the home network scene. One, to get the different machines at home *networked* to allow them to work together, and two, allowing the creation and distribution of media from and to the inhabitants of the home and viewed on a multitude of different devices. Researchers and the industry are working on making both of these developments happen.

One angle that has not received much research attention is the usable security on the home network, and more specifically, *user centred access control on the home network*<sup>1</sup>. Possibly this is because researchers are busy envisioning what a home network really is and what the users of one are supposed to use it for. What is known is that many home network product designers, and notably security designers, do not assume realistic uses for their products – security features of products for domestic use are planned as the family using it consisted of a single inhabitant [Eli02] or that an expert could be available by a phone call [GEND05]. The mindset generally tends to be “first we make it work, then we make it secure, or usable.” (after which the promise is forgotten and in its stead, more features are introduced in the 2.0 release – making security and usability even tougher to add to the project). Access control is either non-existent or where applied, mimics the ways from

---

<sup>1</sup>By this, we are borrowing from the term ‘User centred security’, referring to “security models, mechanisms, systems and software that have usability as a primary motivation or goal.” [ZS96]

business computer network security.

This is where this thesis comes in, trying to synthesize information from researchers, security experts but most importantly from the users themselves, into examining the access control needs and possibilities for the home network. The researchers' input comes from their articles and publications, security experts have voiced their opinions in one-to-one talks and interviews with the author and the actual end users have been heard in on-site family interviews conducted and analyzed with methods and techniques borrowed from Contextual Design.

It is beneficial to examine the security needs of home networks now, as home networks still are an emerging technology. There are challenges to be addressed before home networks go mainstream. One such challenge is allowing controlled access to services on the home network for people outside the home, either using a traditional, largely unstructured client-server model, through peer-to-peer solutions, or as persistent home-to-home connections.

While there are a host of possible threats on the home network's security of technical nature and posed by malicious users, there is both literature [Sta02b] and our research to suggest that access control on the home network is not a very critical issue between legitimate users of the home. Inside the home access control can largely be governed by the social barrier.

### **InHoNets connection**

The research and results in this thesis are connected with the Interconnected Broadband Home Networks project (InHoNets) [inh06], a Helsinki University of Technology (HUT) / Tampere University of Technology (TUT) joint project to research the possibilities and implications on home networks with the focus on interconnecting homes. The project is financed by TEKES and industrial partners. The on-site interviews and their analyzes which are used as material for this thesis, as well as the project in general, were performed by project researchers, including the author.

### **Organization of this thesis**

This thesis is organized as follows: Chapter 2 presents the problem statement and requirements for this thesis, wherein it is presented why usable access control is of any importance in a home network environment and what challenges lie between the current situation and a secure and usable home net. Following are two background chapters: Chapter 3 presents the home network environment, its users and stakeholders, including the external ones, and rounds off by presenting some of the relevant protocols that are of significance on a home network. Of particular practical value is the UPnP<sup>TM</sup>

protocol, with which it is possible to tie together much of the technology at home. Chapter 4 exists as a primer to security, usability, user centred design and Contextual Design, and the merging of these fields into Usable security or HCIsec. There is a natural bias towards what areas of security, usability and HCIsec is topical to the home network ecosystem. Chapter 5 examines previous work in home network security and related research that could be applied to enhance home network security and user-friendly access control.

Chapter 6 describes the user interviews done for this thesis, how they were done, and what was learnt about the home environment and the inhabitants' attitudes towards security, access control and technology in general. It also presents the expert interviews, along with an analysis of them. Further, the chapter has a discussion on what access control solutions would be applicable in a domestic setting. In chapter 7, the research results and the suitability of the methods used in this thesis are analyzed. It is stated that while the methods weren't used to a great depth, they helped the researchers understand both the home network and their users better, as well as the methodology itself. The thesis is concluded in chapter 8.



## Chapter 2

# Problem statement and criteria

The main objectives for this thesis are the following: to find out what kind of access control do users *want* on their home networks, what kind of access control do the users *need* on their home networks, and whether already means exist to *solve* these problems.

While doing this work, the following areas are also surveyed: What is a home network, what are the assets there to protect, and the threats to protect it from, and why the need to protect anything. We look at the needs and requirements of different stakeholders of home network usage, with the main focus given the inhabitants of the home, the users of the home network.

To find out what access control users want, we need to identify our users and their needs. This is done through a form of user interview and analysis borrowed from Contextual design. Since we don't expect to find very developed home networks deployed, we investigate the users' attitudes towards security, privacy and access control, and what kind of devices and media – digital and non-digital, networked and non-networked – people currently have in their homes. We then apply this information on what we know about digital home networks, current and near-future. To enable this, we also need to have a view on what a home network is and what services it provides – or rather, will provide.

Users may not be able to identify all demands and requirements for access control. Much of their understanding of access control is implicit and unarticulated, and they are not aware of the risks brought by the ubiquitous computing environment the home network of the future will be. For that, we need to investigate published literature on security and access control requirements for computer networks in general and those of pervasive computing, and apply the information on a home network context. We also need

to take a look at the problem from the perspective of the service providers, which includes requirements set by legislation. Security expert information is taken from literature and interviews with security experts.

The user requirements and the ones from the security experts need to be analyzed to form a set of consolidated access control requirements. This should show if there is a discrepancy between the users' requirements and those of the security experts, and if so, whether the requirements are conflicting or complementing.

Finally, we ponder whether there are ways to provide solutions to the access control requirements in a usable manner. For this, we use knowledge from the Usable security (HCIsec) field.

## 2.1 Problem statement

Since it cannot be assumed that each home will be inhabited by an expert in security, or even in technology in general, it is paramount that home network security is handled in a usable way. Even the user who does not have an extensive understanding of the underlying intricacies of security must be allowed to use the system in a secure manner and be able to feel that she controls the environment in a secure way. This sets considerable challenges for the design of the home network, which should be of considerable interest to the industry trying to push this technology; for home network technology to be adopted, the home network must be easy to deploy, maintain and use.<sup>1</sup> Current home computer networks are very complex to understand and maintain, much due to the fact that they do not support working models and requirements of families [GEND05, Ell02]. This causes stress and frustration with the users, who will only invent workarounds to circumvent what they see as obstacles set by the system. As more features and services are added to the network without properly addressing these issues, the complexity of it all may become an inhibiting factor for the deployment and acceptance of home network technology.

There are many challenges that need to be tackled. How should access control be presented in an understandable way to a user who most of all just wants to “use” the system?<sup>2</sup> What security models reflect the home users' needs? What means for access control are practical or even acceptable in a home setting, and more so, what means are unacceptable? Do the security

---

<sup>1</sup>These are of course not the only requirements; for a product to succeed, it has not only to provide a good user experience, but also be built on solid technology, and it needs to be efficiently marketed [Nor98]. The cost is a factor, as is timing, design, image and a multitude of other things. All these points however are firmly outside the scope of this thesis.

<sup>2</sup>This very question is being tackled by the author's instructor in her licensiate thesis

demands of users, experts and providers complement or conflict each other? More so, do the access control requirements of different stakeholders – the home network owner, the network service provider, the content creator and distributor, the legislator, etc.– conflict? If they conflict, whose demand should be heeded? How much security can be achieved using access control, and how much access control is really needed in a home environment? People do seem to understand security in other contexts – they understand the importance of locking their doors and keeping their bank cards’ PIN codes to themselves – is it just the security vocabulary, in a broad sense of the word, that is wrong?

There are also many technical challenges for a home network of distributed nature. How shall devices trust each other when there is inherently no top authority on the network? How shall devices be added or removed from the network, especially if they aren’t fully compatible? When separately written programs and appliances are composed so that they may co-operate, they may instead destructively interfere in unanticipated ways [Mil06]. Certainly there need to be concepts borrowed from a network of centralized architecture – service discovery, service registry, security policy service – so that all data about other devices and their access rights need not be contained in each node and distributed through multicast mechanisms. The challenge for the network architects is to provide an architecture that is not dependent on central control, yet supports a fluid user experience. Within the scope of this thesis, some of these questions will be addressed while others have to be left for future work.

## 2.2 Scope

In this thesis, we examine the access control needs and requirements for the home network. There is a bias on on entertainment use – A/V networks, computer use not for work or studies, including communication, and other leisure activities – within single homes. Entertainment is supposedly the most immediate candidate for becoming part of the home network. It is one activity that end users have both an acquaintance with and may be interested in having developed.<sup>3</sup> Interconnected home networks and home automation networks are mentioned briefly, as well as other areas of security than access control. From a cultural standpoint, we are examining the Finnish family.

The choice of focus stems from the connection with the Interconnected

---

<sup>3</sup>The first application users generally want from a ‘smart home’ is remote controlled lighting. This is an affordable, easy to understand application which has little dependencies on other home network fields. Of course, lighting should seldom need to be explicitly controlled at all – it should react to the presence and activities or intent of the user.

Broadband Home Networks (InHoNets [inh06]) research project, in which the author participated. Since the work is based in Finland and involved interviews, the Finnish family was a natural group of users to focus on.

While relevant to home networks and network security, many issues have due to scoping deliberately but regrettably been omitted from this thesis. Such issues are: Trust models and trust management in a ubiquitous network; a deeper inspection of the possibilities, possible extensions and novel applications of UPnP, as well as security implications of the protocol; authorization infrastructures for ubiquitous/nomadic computing [ZK02], and other protocols and architectures usable in a home environment, such as CAN (Controller Area Network, a bus technology initially designed for automotive and industrial applications) [Rob91], Jini<sup>4</sup>, Bluetooth<sup>5</sup> – possibly spiced with Wibree technology for peanut devices – and Zigbee<sup>6</sup> (a protocol suite for low power digital radios, designed to be cheaper and simpler in design than Bluetooth and as such suitable for peanut device use).

---

<sup>4</sup><http://www.JINI.org>

<sup>5</sup><http://www.bluetooth.com> (information site) and <http://www.bluetooth.org> (member site)

<sup>6</sup><http://www.zigbee.com>

## Chapter 3

# The home network

*“The future is already here.  
It’s not just evenly distributed yet.”*  
– William Gibson [Gib93]

### 3.1 What is a home network?

There are several different, often complementing views on what exactly constitutes a home network. The view depends on the field of profession, expertise or interest of the viewer. Computer enthusiasts may consider a home network the actual networking technology used to connect home computers with each other or the Internet. Computer users may include the computers themselves into the equation. Grinter et al [GEND05] consider the home network to be all the computing elements plus the audio/visual (A/V) devices installed in the home, thus noting that the home network consists of the two subnetworks *home computer network* and *the A/V network*. Others have also called for the convergence of the computer and entertainment networks into a home media network [BG02]. The Digital Living Network Alliance [All04] considers consumer electronics, with an emphasis on A/V equipment (‘brown goods’<sup>1</sup>), computers<sup>2</sup> and mobile technology to be parts of the home network, and strives to make this into a single interoperating network. Other researchers include home automation and ‘white goods’ (washing machines, refrigerators, HVAC etc.) a part of the home network [KLKY02], that work in tandem with RFID tags embedded in the clothes, food packaging etc., but what all these visions have in common are that they

---

<sup>1</sup>This term stems from the times when any credible piece of audiovisual equipment would have an aluminum rim and an inset of faux-wood laminate or, in case of film cameras or projectors, faux-leather.

<sup>2</sup>“Beige goods”?

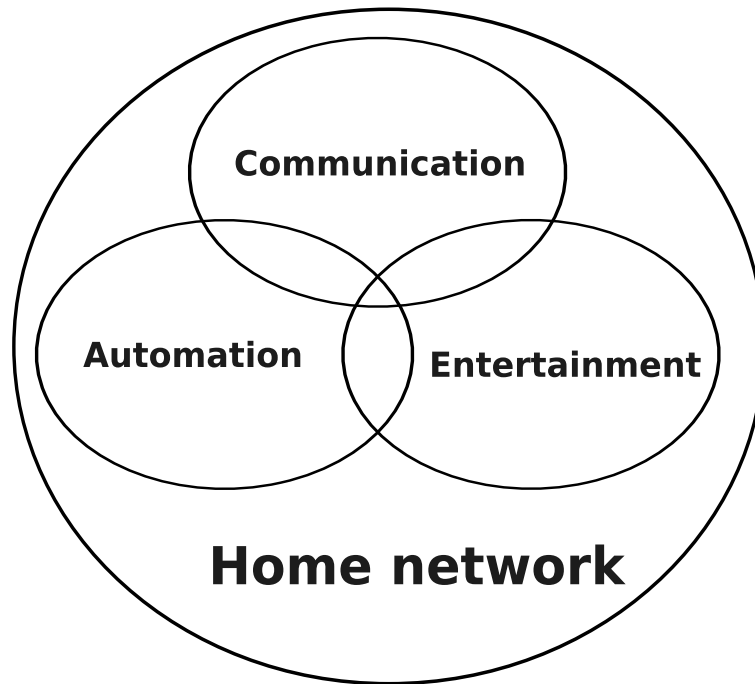


Figure 3.1: The home network environment

are very device-centric.

In this thesis, the home network signifies the interconnected collection of all the devices, technology and services that together facilitate *digital living*. This basically includes any device, appliance or ‘thing’ in the home currently or fittable with a microchip (including its users!) and capable of communicating over a network, now or in a near future, and the ‘things’ and services outside the home, available to the home. This system is examined from a functional perspective; to the user, the network ‘fabric’, the technology and machinery itself, is of secondary importance compared to what services and possibilities it offers to the inhabitants of the home.

Figure 3.1 makes a high-level division of the home network, from a functional perspective, into communication, entertainment and automation. Here, communication part includes phone calls, email and blogging as well as organizational tools such as distributed shopping lists, shared calendars, and reference tools such as wikis. The entertainment part includes enjoying music, pictures (photos, videos, art), games, hobbies and sporting activities. Finally, the automation bit includes traditional aspects such as lighting and climate control as well as security, and most of all, the orchestration and coordination of all the devices working together. The picture omits supporting functions such as the administration (maintenance) of the home network.

|                         |  |
|-------------------------|--|
| 1. Mainframe computing  | One computer serves many users                             |
| 2. Personal computing   | Each user has their own computer                           |
| 3. Personal networks    | Each user has many computing devices, which can co-operate |
| 4. Ubiquitous computing | Many users share the services and power of many computers  |

Table 3.1: The four waves of computing

The development of the home network can be expected to happen in two phases. In the first one, which is already happening, home computing is merging with A/V technology and communication is merging with home computing. Already, “computerless” media players, cameras and Skype- or Windows Messenger-telephones are starting to appear on the market. Control of lights takes a step towards user-centredness by moving control buttons from the walls (where the users *need to be* according to the house-centred view) to remote controls and movement sensors (where the users *are*). Security systems enter the homes and classic thermometers get replaced by wireless weather stations. Few of the devices are connected together as one network.

The next step of the home network is one in the direction of ubiquitous computing environment (see table 3.1). Computers are everywhere, but they don’t look like computers. Instead, they are built into all the things within the house. These devices are connected to a common home network and they can co-operate and they can share knowlege and resources. As an example, a small ubicomp network transaction would be the camera which takes the picture, the GPS which provides location metadata, the cross-reference mechanisms which provides more metadata such as the user’s calendar information, which Bluetooth devices were in vicinity when the picture was taken, and by using image recognition, who or what may be in the picture, and finally the storage mechanisms which put the picture on the family server (which takes an off-site backup of the picture when convenient). The camera itself need only be able to take the picture and call upon the other services.

The ultimate goal of the father of ubiquitous computing Mark Weiser is one where where computing power and digital communications are inexpensive commodities that are embedded in everyday devices and communicating with each other over ad-hoc wireless networks. Minuscule computing devices exist in such a hive-like multitude that computation itself exists as a part of the environment [Wei99]. This means that there will be “a computer in everything”, albeit probably a small and specialized one in most cases. On the other hand, not all devices need to be able, or indeed should have to

handle any task. When the proverbial Internet Toaster needs to know the correct settings for a frozen Karelian pastry, it can consult a data bank that knows. And if the toaster breaks, it can send a notification about this to other devices (and ultimately, its owner) that should be interested. If this sounds like outlandish science fiction to the reader standing at this dawn of the third era of computing, he only needs to be reminded (or informed) that in 1918, Sears Roebuck company sold a “home electric motor” with attachments making it a fan, an egg beater, a sewing machine and a vacuum cleaner [Nor98]<sup>3</sup>.

It may be unrealistic to envision a fully networked home where “everything is connected”. Indeed, there may not be any reason to connect everything in the home, but having the possibility to connect formerly analogue things like sofas and pot plants opens up a whole lot of interesting possibilities.<sup>4</sup>

A home network will need both wired and wireless technologies, depending on application; some applications will be more convenient or flexible when built without wires, other applications may be cheaper, more secure or more reliable when built with cabling. Ultimately, most of the data wires will go away, but it will take some time before we can deliver electricity wirelessly (or have good enough batteries).

The infrastructure to allow for ‘first phase’ home networks is starting to appear in new buildings in the form of structured wiring. There are three levels of structured wiring: power cables are drawn in a star topology (or tree layout in advanced configurations) to allow centralized controlling of electric appliances; signal cabling, usually twisted pair cabling (e.g. UTP-CAT5e or CAT6), also in a hub or tree topology, for computer networking – both as a delivery medium to network endpoints and to work as backline to the wireless infrastructure – wired telephones, television antenna RF signal transmission, distribution of picture and sound in analogue or digital format, home automation control, security applications and anything which can be transferred over the wire on signal strength voltage; and a bus for home control devices like light switches and sensors. As Ethernet only occupies two of the four pairs of a CAT5 cable, the remaining two pairs can be used to transfer Firewire data [MNH<sup>+</sup>01] or even analogue audio as soon as the signal has been converted to a balanced line level audio signal.<sup>5</sup> Added

---

<sup>3</sup>Prior art: David Kline, “The Embedded Internet”, Wired 4.10, October 1996

<sup>4</sup>A sofa could take part in informing the owner that his wallet is within the sofa, and a pot plant’s pot could signal that it’s about time to water the plant... (the caretaker of the plant may still enjoy the watering the plant, so there may not be a need to automate that part)

<sup>5</sup>It is unknown to the author whether twisted pair UTP-CAT5e is a sufficient medium for analogue signal transmission. The published research in this field is scarce indeed. The cable should have no problem transporting an audio frequency signal – CAT5 is specified to transport a 100 MHz signal up to 100 metres on the cable –, but the signal would need

to this, a home could have cabling embedded in the walls for video, line level audio and speaker level signals, though this is most likely to appear within a single room (or between the engine room and cinema room in more advanced configurations). The bus is not strictly necessary but a significant cost saver, as many devices can be connected to the same bus and the control logic doesn't need a separate input for every single control node (i.e. an addressable 'thing' on the network).

Structured wiring was found to be a standard feature on the Finnish Housing Fair in 2006. In contrast, most older houses are not equipped with structured cabling and it can be a considerable hassle to retrofit structured cabling into existing housing. In these buildings, signal transmission can be done wirelessly or over existing telephone, antenna or electric cables, but control over electric appliances such as lights must be done by the device, and not in a centralized fashion. Of course when home network technology advances, power control of appliances more complicated than lightbulbs can be handled at the device. The underlying protocol just needs a request to switch the device between powered and standby modes.

What still is missing from the home network are easily deployable, networking and interoperating devices. Networked media players and WLAN-equipped digital cameras have started to appear on the market, but they build on an existing computer network infrastructure to work. Computers and gaming consoles can, with a little expertise, be configured into media players and many 2007-generation consoles already have this capability built in or available as an upgrade. DVD players are capable of showing digital pictures from memory cards, but few major brand network-capable players exist.

Home automation, computers and A/V equipment are each living on their own, albeit networked island. While hobbyists and enthusiasts have managed to brew their own home networks, home networking technology doesn't yet exist for the typical consumer in the extent implied in this thesis. Therefore it is more appropriate to talk about the *networks* of the home, as suggested in Figure 3.1 on page 10: the computer network, the telecommunications network, the A/V network (lacking an all-connected home network, the audiovisual devices connected together), a possible home automation network and so on, and the extensions of the home networks: the vehicle network (radio, phone, navigation system) and the personal environment (mobile phone, PDA, laptop, and in future visions, authentication tokens). The vision is to interconnect all these networks, devices and *services* and make them work together.

Indeed, the future home network will probably not be a single network op-

---

to be converted into a balanced one, and the 100  $\Omega$  impedance of the UTP cable matched to the rest of the system

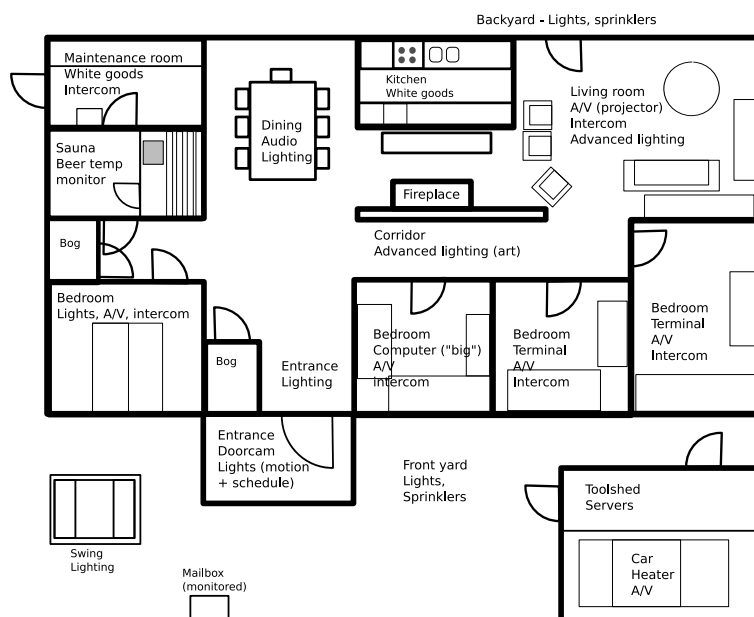


Figure 3.2: A somewhat fictional home network plan

erating with the same networking technology. Different applications will benefit from different underlying technologies designed for the appropriate characteristics – i.e. Ethernet for computers, HAVi for A/V equipment and some bus technology for home automation controllers – operating through a *virtual overlay network* with gateway machines between the networks to allow nodes on the different subnets to interact using a common protocol [NSA02].<sup>6</sup> Protocols already exist for the orchestration of different parts the home network (e.g. UPnP; see § 3.6).

The home network need not be confined to the physical borders of the home. The car and the summer cottage certainly count as a family's extended home and the home network may seamlessly use external services for media storage, access control and so forth. Also, an inhabitant may want to access services the home offers while away from home, such as the media library, a shared shopping list or calendar, security and monitoring, video programming or sauna control (Figure 3.2).

The home computer network is technically not very different from the corresponding one in a business environment but there are many practical issues that set them apart. The home network is dynamic: machines and users may be come and go with little warning. Devices may be turned off or out

<sup>6</sup>Technically, this would constitute a *home internet*

of reach. The home network is heterogeneous, meaning there will be a lot of different appliances on it. There may be a large number of machines per user compared to a business network. The home network is expected to perform a wide array of “tricks”, many of business network complexity, yet work nearly without any training of the users. Most significantly, there may not be a user with very high technical know-how on the home network, and neither should there need to be. The home network is an environment oriented for consumer use. For the home network to gain acceptance, it must not be much more complicated to operate than current stand alone home technology. In fact, it should be a lot simpler.

Not only the technology, but also the users of the home network are quite different from those of office network.<sup>7</sup> Foremost, the users’ goals are different; they just want “to use” the network and its services [SG02]. While every business network should have expert personnel on IT management and support, no such assumption can be made for the home network. As an answer to this, some domestic devices (notably, computers and digital TV receivers) come with auto-updating features. A home network typically has less than ten internal users. By offering services to users outside the home network, either publicly or using defined home-to-home connection, the home network user count naturally goes up.

One characteristic that may have serious implications when it comes to automatic identification of users (e.g. with biometric identification) is that one device may have several simultaneous users, or several users nearby the device but the identification system cannot be sure which user is actually using the device. It is unknown whether there actually *needs* to be an identification system supporting multiple concurrent users; in practise the user with most access rights should be able to handle the access control practicalities. Technically, access control in this scenario may be solved with a model based on Lattice-Based Access Control (LBAC) [Den76]. The effective access rights of the group may, depending on the application, be the union of all access rights, or the intersection.

One architectural assumption of the InNoNets research project is that there might not be a central controlling server on the home network. The assumption is based on the fact that users are likely to be more interested in purchasing devices “that actually do something” rather than investing in the invisible infrastructure. This creates a challenge both for device management and access control, as there is no authoratary device on the net which all other devices can trust. Some building blocks for a decentralized device management and access control are presented in this thesis (§ 3.6, §

---

<sup>7</sup>There should be an interesting discussion regarding the fact that many of the home network users actually are office network users when at work in an office environment. How do their actions and expectations differ between these two environments?

5.4, §7.3).

While home computer security is a well publicized subject, not much academic research has been done specifically for security on home networks. Some of the security research, methodologies and best practises for business computer networks are relevant for home use, but not all of it is very well suited for the home network due to the home network's characteristics.

### 3.2 Terminology soup

Home networks and related technology have been described with many different names. A *Local Area Network* (LAN) is – depending on the context – *either* a network of computers, printers, servers, terminals and network devices, *or* the networking technology needed to connect the computers, printers and devices together. The LAN is usually connected through a gateway to a parent network, or a *Wide Area Network* (WAN), or, depending on scope and installation, through multiple gateways to other LANs, or an intermediate Campus Area Network. A *Home Area Network* (HAN) is a domestic LAN, contained within a user's home that connects a person's digital devices, from multiple computers and their peripheral devices to telephones, VCRs, televisions, video games, home security systems, “smart” appliances, fax machines and other digital devices that are wired into the network. A *Home Network*, in the context of this thesis, is a HAN encompassing all networkable digital devices in a home, including home automation, security, ‘brown’ and ‘white’ goods. A *Smart Home* is a home with a home network, which can react to its inhabitants' needs without explicit user input.

In *Pervasive Computing* and *Ubiquitous computing* [Wei99], computers exist distributed in such large quantity and small size, essentially everywhere, that they effectively disappear. The computers “understand” their users and the users' context and computing itself happens in the ubiquitous environment, meaning that programs and processing can happen in any device convenient that the user has access to. The rationale here is that instead of monolithic ‘kitchen-sink’ devices packed with every conceivable feature, there would be small specialized deviced ones capable of ad-hoc networking. The concept of Ubiquitous computing has also been known as “things that think”<sup>8</sup>, or ‘everyware’. [Gre06]. Indeed, ubiquitous computing can be divided into pervasive computing, nomadic computing and ad-hoc computing [Sta02b], but this is already outside the scope of this thesis.

---

<sup>8</sup><http://ttd.media.mit.edu/vision/vision.html>

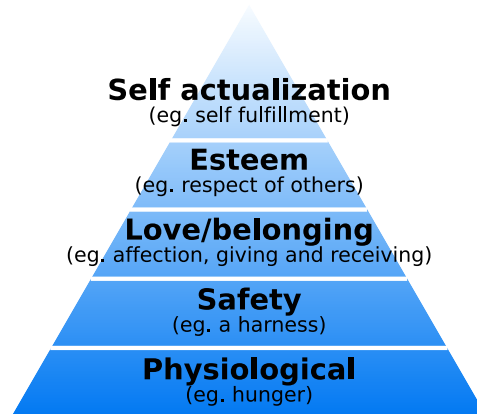


Figure 3.3: Abraham Maslow's Hierarchy of needs

### 3.3 Requirements for a home network

The human needs can, in order of potency, be described as physiological, safety, love/belonging, esteem and actualization [Mas43], as identified by Abraham Maslow (see Figure 3.3).<sup>9</sup> Milton Rokeach further identified 36 human values that describe desired modes of conduct (capable, courageous, independent, responsible...) and desired end-values (a comfortable life, a world at peace, self-respect...) [Rok73]. A home should support these needs and values in providing its inhabitants a safe and secure habitat, healthy and convenient to live in. Ideally, a home should provide both entertainment and give its inhabitants a sense of belonging and possibility of self-expression. Users also expect home network technology to save them time and money, relieve them from organizational stress, allow them flexibility, mobility and security, and help elderly or disabled people [RHB03]. With the technical advances, a networked home can help to provide new and better solutions to these requirements through sensors, monitoring and automation.

On a higher level of abstraction, the house should be *aware* of its inhabitants and their desires like an electronic butler, and be able to react in a way that seems predictable to the inhabitants. For this, the home network needs to have sensors, such as ultrasound sensors and “smart floors” to sense who in the house is doing what and where [EKO<sup>+</sup>99]. It also needs to learn from its inhabitants behaviour, using calendar information and artificial intelligence applications.

For the user, it is essential that she can trust the home technology and that it is easy to use. For that, different parts of the home network must be able to work together. Bodies such as the Digital Living Network Alliance

---

<sup>9</sup>It must be noted that some journalists and bloggers in oppressed countries have proven that in some cases, actualization can be a more powerful driver than the need for safety.

(DLNA) are working on the architecture work needed for a seamless integrated home network environment [All04]. The technology, and especially the security technology – the focus of this thesis –, should be as transparent and unobtrusive as possible, and secure the resources with relevant measures. Devices and services should as far as possible be able to configure themselves with near zero user interaction. Resources of lesser security demands can be authenticated to transparently and only scenarios requiring a higher level of assurance or greater security may require users to explicitly interact with the security subsystem [CAMN<sup>+</sup>02]. Security should also react to context, for example when no inhabitant is in the house, the house enters “a higher level of paranoia” [Kar06].

The home network must be built on the concept of transparency (i.e. the network doesn’t unnecessarily hide its doings) and feedback. Users must be able to easily see “what the house is doing” and receive proper feedback on the actions requested (e.g. “Your coffee is being brewed”), since a user cannot often see the effects of his action [RHB03].

After technology, functionality and usability have been accommodated for, an important step remains: As products aimed for the consumer, the home network should be fun to use. Users of productivity tools have a task to accomplish and do not want to be interrupted, but users of home *entertainment* network want just that: to be entertained. The massive market for cell phone ringtones should be sufficient indication that users want fun technology. Fun is the opposite of frustration, a thing usability tries to eliminate. Designers are now beginning to develop theories of engagement through fun-features [Sch04a]. Charlotte Wiberg has created a set of heuristics to evaluate the fun-ness of websites [Wib03, Wib05]. While her *funology heuristics* are clearly most applicable for the transient and exploratory nature of websites, it should be possible to produce similar heuristics for home networks. While not published in a scientific context, Kathy Sierra proposes that the level above usability would be *Flow/Enchantment*, described as “*Does it keep me fully engaged, where the world drops away?*” [Sie07]. This certainly would be a desirable attribute for home entertainment.

Research has shown that ordinary home-dwellers are ready to accept pervasive technology into their homes, as long as it supports their goals and they feel they can be in control of it [RHB03].

### 3.4 Homes and their ‘users’

Solutions that are meant to be built and sold for home use must be built for the users of the home, i.e. the inhabitants. This sets a new challenge for the security designers, who traditionally have created solutions for business use, which is much more straightforward than the domestic setting.

The simplest form of a “home user configuration” is that of the single inhabitant. All the devices are owned and controlled by this person, and there is no need for any access control from the inside of the net. This is the most basic and simplistic home, and the one many security designers assume [Ell02]. The picture does not change significantly in the case of a family of two adults, or families with small children, as long as the adults can agree on the security policy.

On the other side of the spectrum comes the family with teenagers. Teenagers want a degree of independence and privacy and probably also own personal networked devices. Furthermore, they may invite friends over who want to connect their devices to the network.

A special case of interconnected homes are separated families; the children would need access to their data regardless of which home they are in, while the separated parents probably would like to keep their resources inaccessible from each other. Another special case of both interconnected homes and rights management arises with homes inhabited solely by elders which may not be able to fully take care of themselves. A responsible caretaker needs to have considerable access to the elders’ dwelling (to get alerted if there is an acute health problem, for instance), while still providing careful privacy of the home.

To further discourage the notion of the single user home, one only needs to note that a home is generally inhabited by several users with different backgrounds and different requirements [RHB03].

Inhabitants of the home may want to be connected with their peer groups – hobby groups, friends, relatives, the children of separated families – to share data and resources on their network. This calls for interconnecting the home networks, for which there is quite a lot of current research activity ([inh06] etc). A home’s technology should also allow for flexible rights management to allow for occasional visitors such as servicemen or the neighbour who waters the plants to enter.

Our interviews indicated that home users would prefer not to be “bothered by computer security.” Still, people tend to be conscious about their real-life security. Further research should be able to point out reasons for this discrepancy – do people lack the needed parallels between computer security and real life security?

### 3.5 Stakeholders

In the game of home networks, there are many stakeholders with different and even conflicting goals: the users (home inhabitants and their peers), the media publishers and the service providers (ISPs and content providers).

The authorities are scoped out from this discussion as stakeholders: from this point of view, authorities set the rules for the benefit of the other identified stakeholders.

The main stakeholder in this thesis is the inhabitant of the home, whose main goal is to *use* the home network and the features it offers. The network and the devices on it are not of great importance to the user *per se*, rather the services it provides (this may not be entirely true; some users may take pleasure in having the home network machinery itself). Maintaining security is a tertiary goal at best. The home network should be so easy to use for the user that it almost becomes invisible [Wei99].

No inhabitant is an island, and so all of them have peers, organizable into peer groups. A peer group may be the extended family, a group of friends, can be related to different hobbies, and can be very temporary in nature as in the case of a group of people organizing a stag/hen night for a common friend. Peer groups want to be able to access data, and possibly services, on the home network, in an easy manner. Services should be easily discoverable. In case of more permanent relations, homes may be interconnected for recurring access to the resources. From the inhabitants' point of view, material should be visible in appropriate proportions and secure from unwanted eyes. Interconnecting homes and the related problematics are the focus of research of the InHoNets project [inh06].

A different set of stakeholders are the bodies creating services for the users, and making money from that activity. Service providers can make money either directly through charging for the service they provide, or indirectly through for example advertising.

Media creators, publishers and aggregators – for example film companies, studios and television channels – create and distribute material to play on the entertainment portion of the home network. Media generators are interested in delivering their media in a safe and efficient way and to maximize on their capital. The media is their way of income which is why the distributors want it protected from free use.

Internet service providers (ISPs) provide the user an Internet connection ('bandwidth') and, at an increasing level, services and content. The content can be offered at no extra charge ("for free") or at a premium. Typical services are email and web space and security services. ISPs want Quality of Service on their network and the media, because that's what they sell. From the ISPs point of view (and that of the Legislator), an infected or ill-behaving machine is "problem waste" on the network.

Equipment manufacturers have a stake in the game too and a reputation to protect. If the image of a product suffers due to the product being unreliable or unfit for its purposes, consumers are going to turn to another

brand of products. A product may be deemed unfit because of actual lack in functionality – security, for example – or it may just be too complicated to use. Additionally, users may choose a certain device just on not rational grounds, but also emotional ones.

### 3.6 Backline technology: Universal Plug and Play

UPnP<sup>TM</sup> is a family of protocols for autonomous network setup and management, allowing controllable *devices* and *control points* on a network to communicate and interoperate [UPn03]. UPnP has features needed for a device joining and leaving a network, service discovery and service announcement, and control and interoperation of home network devices, without the need of a central control server. For example, one can have one or more media servers on the network, which will be picked up by the media renderers (players) on the net, which in turn will be picked up by the control devices, all with minimal input from the user. The protocols are extendable and support access control and distributed and delegated security.

UPnP seems a compelling and well defined technology for the home network backline and implementations have already been made both for research and production [NL94]. It can work over any transport medium, so computer network devices can continue operating over Ethernet while light switches, dishwashers and coffee brewers can operate over Zigbee, CAN, Bluetooth or whatever network technology is convenient. UPnP builds on HTTP, XML and SOAP and uses common protocols instead of vendor-specific device drivers (yet allows for vendor-specific extensions). UPnP devices can be implemented on any language and operating system.

While the control of media players is a part of the UPnP technology, transport of media, e.g. from a media server to a media player, is not. However, since UPnP is an expandable – even vendor-expandable – technology, there is no reason why this couldn't be done as a UPnP extension.

UPnP has steps for device *discovery*, *description*, *control*, *event notification* and *presentation*. UPnP's support for access control security is based on Access Control Lists (ACLs) and digital signatures [El103], and managed using Security Consoles through a specified Security Ceremony. The use of UPnP security measures are optional and not yet widely adopted.

## UPnP base operations

When a node joins a network, it will first seek for a DHCP server for an IP address and other network information, such as a hostname<sup>10</sup>. If a DHCP server isn't found, the node assigns itself an IP address using Auto-IP (in short: 1. select an arbitrary address within the 169.254.1.0–169.254.254.244 range, 2. use ARP to check that this address is not in use and repeat from 1. if necessary; 3. send by ARP that the address has been chosen; 4. periodically check for address collision and go to 1. if one happens). A device will then advertise a few of its essential specifics, e.g. its type, UUID, duration until service expires, and an URL to more detailed information, using a number of multicast messages corresponding to each of its embedded devices and services to a standard address and port (239.255.255.250:1900). Similarly, a control point will multicast a discovery message to which devices respond with an advertisement message. Control points can later send description queries to devices they are interested in. When a device or a control point leaves the network, it should (if possible) send out revocation messages by multicast. Devices also periodically re-send the advertisements with a new expiration time if they are still available. To limit network congestion, the time-to-live (TTL) of each multicast IP packet should default to 4 and should be configurable.

The next step in UPnP networking is Description. A control point may request more information about a device and the services it provides using the URL the device advertised above. This information is in XML and generally retrieved over HTTP. A device will carry a *device description* describing its manufacturer information (model name and number, serial number, etc.) and *service descriptions* which is based either on a standard UPnP *device template* or a vendor's extension and includes the *actions* the service responds to (i.e. commands), the arguments for each action, a list of variables which reflect the state of the service at run time and URLs for control, eventing and presentation. A control point can now ask a service to invoke those actions by sending a *control* message to the control URL of the service, and receive feedback on the results of these requests, or ask the value of a service's state variables.

The inverse of control in UPnP is *eventing*: services send information about change of state of any eventable variables to control points that subscribe to event messages (non-eventable variables need to be polled for explicitly). The event notification can also contain encrypted information that can only be opened by authorized Control Points.

---

<sup>10</sup>DHCP has hundreds of options to describe the basic services of a structured network such as the location of a time server, mail and web servers and the location of a boot image for a node which gets its system over the network, such as diskless terminals and the Hauppauge MVP media player.

|                         |  |
|-------------------------|--|
| basic device            | Does not provide any services. Useful for compliance testing and extending by vendor   |
| media renderer          | Audio, video and/or picture "player"   |
| media server            | Server for audio, video and picture content  |
| lighting controls       | DCPs to describe binary light and dimmable light devices, and dimming and switch power services to manipulate them                                     |
| HVAC                    | Fans, temperature sensors and thermostats, needed for climate control  |
| digital security camera | Control a security camera with motion detection  |
| Internet gateway        | Contains DCPs for gateways, (TCP/IP) LAN devices, different kinds of WAN devices, and configuration of all of them                                     |
| WLAN access point       | DCP to access and control network and security aspects of a WLAN access point  |
| printers                | Allows printing and print job control in "basic" and "enhanced" flavours   |
| scanner                 | DCP to control a scanner with a document feeder and a control panel  |
| remote UI               | Describes the server and client devices and services to create and display a user interface which optional support for interaction and device security |
| device security         | Security services for UPnP, including the security console   |
| quality of service      | Describes QoS services that a device may have or control   |

Table 3.2: Standardized UPnP Device Control Protocols (DCPs)

Finally, UPnP services can provide a HTML *presentation* which can include state information and means for the user to control the device.

A set of standardized UPnP device classes have been presented by the UPnP forum as standardized Device Control Protocol (DCP) descriptions. The most relevant ones to this work are listed in table 3.2

### UPnP security

UPnP security uses a combination of Device and Control Point called the *Security Console*. Its purpose is to take over security ownership of the Devices and then to authorize Control Points or other Security Consoles access to Devices the Security Console administers [E1102]. The security orchestration relies on public keys. A new device would report the SHA-1 hash of its public key to the *Security Console* at which the user can verify the reported Security ID key is what it should (e.g. by comparing it with a printed card shipped with the device). The user is then prompted to give

a meaningful name to the device, which the device from then on will be referred as. This ceremony should be quite adaptable to the ones described in § 5.4, which would take out the need to manually compare SHA-1 hashes, a typically error-prone task for humans. Any arbitrary degree of security authorization, including the delegation of a devices access control list (ACL) can also be delegated and shared among other security consoles, which is a useful feature on a distributed network. Delegation of ACLs is a measure both against power cycling of devices and a benefit for devices with such a small memory footprint that it is more efficient to have its ACL elsewhere.

UPnP has been criticized for its flaws in security. These flaws have mostly been related to the implementation of UPnP and are not inherent to the protocol itself. UPnP has also been criticized because the protocol does not have provisioning against denial of service (DoS) attacks, but DoS attacks could just as well be thwarted by the UPnP implementation or a separate subsystem such as a firewall. Security-critical applications such as firewalls must take extra provisioning against malicious use if they are to be configurable by UPnP.<sup>11</sup>

### Alternatives to UPnP

UPnP can be overly heavyweight for ‘peanut devices’ due to the connectivity requirements (typically Ethernet), the amount of processing needed and amount of data associated with the messaging (as introduced by the XML overhead).<sup>12</sup> Since UPnP allows a root device to have any number of services and any number embedded of physical or virtual devices (each with an arbitrary number of services), a whole collection of small appliances (e.g. all lights in a room) can be connected as devices to the UPnP root device [KLKY02].

Near relatives to UPnP are Zeroconf, a set of technologies for device network autoconfiguration and service discovery<sup>13</sup> (called Bonjour on the Macintosh OSX), the Jini and OSGi technologies, rooted in Java, and Microsoft Rally, which is like UPnP with added features for Quality of Service (QoS, for streaming media) and easier configuration of wireless devices.<sup>14</sup> Other related technologies, such as Bluetooth, Salutation, Service location protocol (SLP), Secure service discovery service (SSDS), Centaurus and the unnamed protocol by Burnside et al are omitted from this text, but are disseminated

---

<sup>11</sup>Yaron Goland, one of the designers of the ill-fated Microsoft UPnP stack, comments the security flaws in a blog entry from 2002 at [http://www.goland.org/upnp\\_security\\_flaws](http://www.goland.org/upnp_security_flaws)

<sup>12</sup>But factoring in Moore’s Law, vastly more complex Peanut devices await us

<sup>13</sup>Unlike UPnP, Zeroconf can seek the answer to questions like “who on this network can print?”

<sup>14</sup><http://www.zeroconf.org/>, <http://sun.com/jini/>, <http://www.osgi.org/> and <http://www.microsoft.com/rally> respectively.

in [CGR04].



## Chapter 4

# Usable security

In the days when computers were operated by trained experts in white lab coats, there was little concern for either security or usability in computing. Unfortunately the perception prevailed when computers started to become more commonplace; many of today's problems with network security stem from a time when the Internet was considered a friendly place of researchers and academics and there was little need for network security. Security was glued on to products more as patches to existing processes. Usability and "user friendliness" was a thing measured at the end of the product development stage in usability labs, if at all.

Currently it is known that both security and usability are things that can – and should – be an integral part of product development from the very start. Rules, guidelines, methods and best practises exist for the fields of usability and security so that they can be thought of proactively, but still usability and security are often seen as different camps, and that usability and security as properties are mutually exclusive. The current challenge is to couple the old antagonists usability and security together into what is coined Usable security.

Ka-Ping Yee argues that this setting has historical grounds; if adding security to a product at a late stage, usability suffers and if adding usability to a product at a late stage, security suffers [Yee04]. Yee goes on to demonstrate that if thought of from the start, usability and security can work together to create secure and usable products.

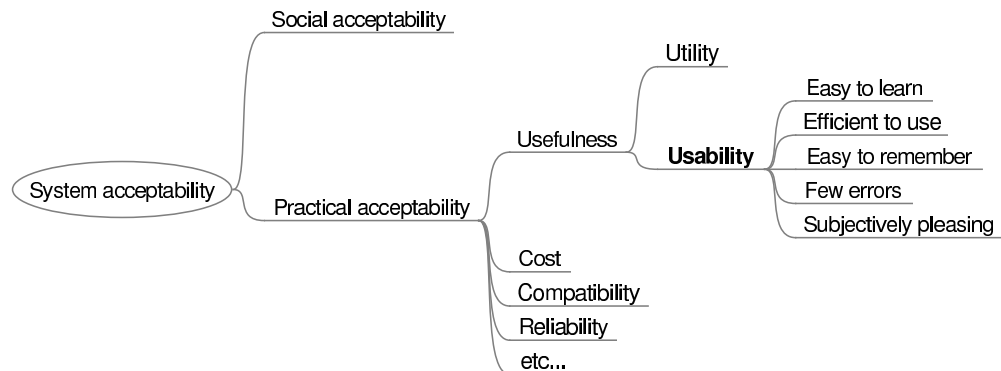


Figure 4.1: The facets of acceptability

## 4.1 Usability and User centred design

*“Don’t make me think.”*

– Steve Krug

A central quality indicator of any thing or process is its *usefulness*. Usefulness is defined as whether the thing can be used to achieve some desired goal, and is the product of its *utility* and *usability* [Gru92]. Utility is what the thing can do, with respect to what it needs to do. Usability is a way to ensure, or at least enable, safe and efficient use of the thing for its intended purpose. This can happen when the product is designed with the users’ psychology and physiology in mind. The International Organization for Standardization defines usability as “the extent to which a product can be used by specified users to achieve specified goals with effectiveness, efficiency and satisfaction in a specified context of use.” [ISO98], or in less grandiose terms, something that *efficiently* gets the job done.<sup>1</sup>

Usability is often used as a synonym for Human-Computer Interaction (HCI), but usability can be related to anything that can be *used*, not just computers. Usability can be used in non-tangible sense, as in “the usability of documentation”.

In addition to being safe and efficient to use, a usable product also needs to be reasonably easy to *learn*. It should be easy to *remember* how one would use it even if it was some time between uses. The product should be designed so that users make *few errors* using it, and when they do, there should be an easy way to recover from the error. Finally, the system should

<sup>1</sup>A close relative to the term usefulness is *user experience* – ‘UE’ or now more commonly ‘UX’ – relating to all aspects of the user’s interactions with the product; how it is perceived, learned, used, and the needs the product fulfills [Nor98].




Figure 4.2: Floor plan of typical (but hypothetical) usability lab

be *pleasant to use*. These characteristics – along with safe and efficient use – were dubbed the *usability heuristics* by Jakob Nielsen [Nie93] and are illustrated, with their parent attributes, in Figure 4.1 (picture adapted from [Nie93]).

In the early days of usability, the tools for creating a usable product were guidelines to achieve the above usability heuristics, and testing. The testing was done as a quality assurance towards the end of the product development cycle and preferably performed in a usability testing lab. A usability test lab would consist of a user space and a researcher space, separated by a one-way mirror (Figure ??, origin unknown). The user space is created to feel as homely and un-lablike as possible for the user. The user and an interviewer sit on one side of the mirror while usability experts occupy the other half of the lab. The user is then given typical tasks to perform with the product tested while the experts record and observe the testing from their side. It is important to tell the user that the usability people are not testing her, but the product she is testing. The results are later disseminated to see where the user had problems completing the given tasks.

It soon became obvious that errors in the design were easier to fix earlier on in the design cycle. Thus, usability testing began using interface mock-ups and paper prototypes, crude representations of the application tested. Even a web site's navigation system could be simulated using a stack of papers with page numbers for hyperlinks. Other developments in usability testing included the Cognitive walkthrough and the Heuristic evaluation.

With these, the usability expert would ask himself “As a user, what would i do now?”, and “How does this comply with the usability heuristics?”, respectively.

While the methods just described are in no way outdated, they are now but a part of the current usability trend, User Centred Design (UCD). The key principles for UCD is to focus on users and their tasks early in the design process by actually involving the users in the design process, to use empirical data – actual test results with actual users – to evaluate the product, and to design iteratively using a design, test and measure, redesign and repeat as necessary cycle [GL85]. According to UCD, the product should be built for the user, thus fulfilling the user’s needs. The input for *what* should be built needs to come from the user, which is a giant leap from the earlier perception that the functionality of the product should be specified by the developer<sup>2</sup>. *How* the thing should be built would still be the task of the developers and usability people.

There are a few notable UCD methods, and the one we shall focus on in this thesis is Contextual Design.

## 4.2 Contextual Design

*“See first, think later, then test. But always see first.  
Otherwise you will only see what you were expecting.  
Most scientists forget that.”*  
– Wonko The Sane [Ada84]

Contextual Design (CD) [BH97] is an approach to designing products and processes that the customer needs, based on the customers’ needs. The process is primarily intended to create new solutions for existing problems, which is in slight contrast with the goals of this thesis: to create new solutions to a new problem. The creators of Contextual Design however do state that CD also can be used for a scenario like ours; people are probably doing similar things in other settings or using different technology to achieve similar goals. It is up to the designers to understand what to apply to the new scenario.

The basis for Contextual Design is in understanding the user and her needs, as should be the basis for any kind of design. While understanding the user sounds like a painfully obvious starting point for any development, specifications for what is to be built in reality often originate from the management,

---

<sup>2</sup>Or, as often has been the case, the marketing

the marketing department, focus groups, market research, key users, or simply as an educated hunch from the developers themselves. The information on what the users need is most readily available from the users themselves. The challenge, for which the CD process provides a tool, is to fish out this information and transform it into product requirements.

The CD process consists of the following steps: Contextual Inquiry, Work Modelling, Consolidation, Work redesign, User environment design, testing with the customer (i.e. Prototyping) and Implementation design.

**Contextual Inquiry** The Contextual Design process starts with the Contextual Inquiry (also: Contextual *Enquiry*). This is a set of well-planned master-apprentice type interviews with the actual end user of the product, at the site of the end user. Basically the idea is to watch the user do his work and to ask questions like “why are you doing this?” or say “show me!” to dig out the silent, implicit information about the user’s work. The goal of the Contextual Inquiry is to get into the head of the user, to see the work in its environment through the eyes of the user and to create a *shared interpretation* of the work between the user and the interviewer. For the user, the work has become so habitual that they don’t even think about the steps involved and have a difficulty articulating both what they did, how they did it, and why when not actually doing it (i.e. in the very context of work). Using contextual inquiry techniques, this information can be made explicit.

**Interpretation and modelling** The information generated by the interviewers is disseminated at an Interpretation session with a cross-functional team, including both developers and business people. Having people with views from all sides present at the interpretation session serves two distinct purposes. The end user’s work is examined from all relevant perspectives so that all members gain insight of the situation from all sides, and the whole design team gain a shared understanding on what should be built, and why. The developers can now focus on solving explicitly the issues the customer needs to have solved and not rely on their own guessing.

The interpretation is aided through the use of different *work models*, which illustrate the whole work process but from different angles. This is an important communications tool to describe the user’s way of working, and especially so when the work is complex and unfamiliar to team members

**Consolidation** The modelled interview data from all the interviewed customers are brought together so that the design team can see common patterns, as well as individual variations. This consolidated information is used to create a solution that will cater all the users’ needs, and that will work

well within the whole receiving organization, or at a higher level, to create a solution for a whole customer population – the whole market of multiple customers. The work is done using an Affinity diagram or Affinity wall, essentially a set of Stick-it<sup>TM</sup> notes in different colours organized by theme with series of connecting lines. A large affinity diagram can occupy a whole room, and provides an excellent area for the further steps of Contextual design, since the designers are effectively surrounded by the user data.

**Work redesign** When the users' work models have been charted, the design team can discuss new ways to solve the existing problems. The team should now have a solid foundation of the users' work and are able to innovate improved ways how the work could be done. The team focuses on how to improve the work practise using *storyboards*, the CD equivalent of user requirements, use cases or user stories. These text and graphic illustrations, not quite unlike those used in cinematography, illustrate what a given part of application is going to solve and how on a functional level.

**User environment design** This step is the first one which actually deals with the system the designers are going to build. Designers create a User environment design, a 'floor plan', so that the new system will have the appropriate function and structure to support a natural flow of work. The floor plan shows the parts of the system, what features and functionality exist in each part, how they support the user's work, and how they relate to each other for the user's point of view. It is a tool to help dividing the system design work among developers while maintaining a coherent system for the user.

**Prototyping** The floor plan created in the previous step serves as a tool not only for the system designers but also the user interface (UI) designers. UI designers create user interface mock-ups, paper prototypes, which are iteratively tested and re-designed together with the customer before any of the user interface is committed into code, and where even radical changes to the user interface are comparatively inexpensive to make.

**Implementation** The designed solutions can now be prioritized for an incremental roll-out and translated into code.

There is also a more agile variant of Contextual Design called Rapid Contextual Design [HWW05], which is designed to be easier to adapt to a customer's – or developer's – existing design process and to allow certain design steps to be left out.

In the research for this thesis, we use inquiry and analysis (interpretation, and consolidation) steps from contextual design. The design phase is touched, but mostly left to a possible Licentiate thesis. The outcome of the inquiry forms the basis on use cases on which the InHoNets project can work.

### 4.3 Security basics

*“Staff security is not a luxury. It is not an option. It is a necessity and an essential part of the cost of doing business.”*

– UN Secretary-General Kofi Annan

Security is the basic service that enables the user to do what she wants to do without being disturbed or disrupted by “an uninvited guest”. A good start for security management is to document a security policy. This is a high level statement of what is allowed on the network and what is not, what the acceptable risk is and how security is to be upheld. In a home context, this could include that only authorized processes, devices and users are allowed to access services on the local network, the perimeter security system shall be armed when no inhabitant is around, and no surfing for the kids after bedtime. A start like this should be comfortable even for an inhabitant of lesser technological expertise, especially if the home network products do help the owner with sensible defaults. Security decisions of a more technical and precise kind can then be made based on the security policy.

The field of security is usually divided into *confidentiality, integrity and availability* (‘CIA’), described thus by The Computer Emergency Response Team (CERT): [CER01]<sup>3</sup>

**Confidentiality:** information should be available only to those who rightfully have access to it,

**Integrity:** information should be modified only by those who are authorized to do so,

**Availability:** information should be accessible to those who need it when they need it.

The above definitions can well be generalized to cover all resources of the

---

<sup>3</sup>Oddly enough, the CIA list translates rather well to the control of who can *read, write* and *execute* files. Again, while the author hasn’t seen this parallel in writing anywhere before, he acknowledges he may be re-inventing the wheel.

home network, not just information. For instance, the Internet connection should be available to those authorized, when so needed.

Confidentiality means that information remains undisclosed to unauthorized parties. Confidentiality is (formally) a property of the data, and its purpose is to state who are allowed to discover the contents of that data. It is upheld by security mechanisms such as devices on the net providing access control – so that an unauthorized party cannot get to the data – or by encryption – so that the attacker cannot understand the message – or both.

Integrity is also a property of the data. Integrity means that the information in a message is correct, that it has not been modified by an unauthorized party (and that it was correct to begin with), or more practically, that the contents of the message cannot be altered without the authorized principals noticing. This property would be important for example for a clock service on the network or some data files that become unusable if their integrity was compromised (i.e. were “corrupted”). A popular way of upholding integrity is adding check values to the data such as cyclical redundancy checks (CRCs) or hashes. These values can be publicly readable (and thus, alterable), or modified to be secret to all but the authorized principals. Two common “secret” check value mechanisms are the message authentication code (MAC), which conceptually is the hash of the message and a shared secret, and digital signatures which basically is MAC codes realized with public key cryptography. Table 4.1 compares the different integrity protecting mechanisms presented above.

|                  | Who can generate it? | Who can verify it? |
|------------------|----------------------|--------------------|
| <b>Hash</b>      | Everyone             | Everyone           |
| <b>MAC</b>       | Holders of secret    | Holders of secret  |
| <b>Signature</b> | Holder of secret     | Everyone           |

Table 4.1: Comparison of integrity protecting mechanisms (from [Sta02a])

On home networks, the CIA list needs to be inspected in a slightly different light than originally intended for information security. Within a home network context, and especially from the user’s point of view, a central property for personal perceived security is *privacy*. From the user’s perspective, one could say that privacy is the outcome of confidentiality and integrity.

Still for the home user, *availability* is the most important property of the CIA trio; if a device or a file on the home network is not available, it is of no use for the user [SA99]. A device can be knocked out or hindered with a denial of service attack (DoS) or in the case of a small, battery powered node, a sleep deprivation attack (essentially, a slow DoS attack) which would drain the node of all its power. A file can be deleted, moved or renamed by mistake or by a malicious entity, rendering it unavailable.

Access control is a way to ensure both privacy, integrity and availability.

There is no perfectly secure system.<sup>4</sup> Any real life system has *vulnerabilities*, i.e. weaknesses in the system, be they technical or induced by users. There are *technical vulnerabilities* such as buffer overflows, protocol timing attacks, message replays and so on, and *social vulnerabilities* such as users forgetting to backup files or locking their computers, or users revealing their passwords through social engineering tricks [FSH03].

Vulnerabilities afford *threats*, “things that can go wrong”. A threat becomes a *security breach* through a *failure*, which happens through the combination of a *latent failure* (weakness in the system) and an *active failure* (slips, lapses, mistakes and violations). A *risk* is the expected likelihood of the threat multiplied by the damage or *loss* it brings if it would realize. Security management is finding and maintaining a cost effective balance between what should be protected and to what a degree. It is not worth putting more effort into countermeasures than what the loss would cost, measured in time, money, reputation or other damage, if it occurred.

Computer security is just a part of information security. There is only so much you can do to solve security using technical means. Security begins with the people, who need to be sufficiently informed about security issues. For example. one way of causing damage to another’s network is to apply social engineering tricks [MS02]. If the home network security is based on an understandable policy, it will be a much more comprehensible task for the inhabitants to keep their home secure.

## 4.4 Access control

Access control is the activity of permitting allowed users, devices and processes certain *access operations* to given resources and processes on the network while keeping the unwanted ones away. On a computer system, typical access operations are read, write, execute, delete, change permissions. On a print system, relevant access operations would be print and manipulate the print queue. On a home network, different devices would have access operations relevant to the device itself; read the temperature, record a TV show, unlock a door, change the lighting or HVAC settings.

Access control is divided into *identification*, *authentication*, *authorization* and *audit* (also *accounting*, *accountability* or *traceability*). Users, devices and processes capable of doing something on the net are called as *actors*

---

<sup>4</sup>It has been argued that a disconnected computer encased in concrete and thrown in the Mariana Trench is a perfectly secure system. It has also been argued that from a user perspective, this isn’t much of a system, let alone a very usable one.

or *subjects* or – especially in security texts – *principals*.<sup>5</sup> The terminology varies between authors and a discipline; in this text, the term ‘actor’ will be used. A *resource* is something on the net that can be acted upon. A *device* or *appliance* is the physical manifestation of a resource, while a *service* is something intangible a resource offers or can be requested to do.

Identification is the process of establishing who the actor is to another resource on the network (user: “I am Bob”). Authentication is verifying this identity (machine: “Prove it”). The identification and authentication may also be coupled with a request to do something on a device on the network. The system would then decide whether to authorize the actor to perform the requested action (machine: “ ‘lo Bob. I’ll grant you these rights”). While authorization usually is arbitrated on actor credentials, an access decision can be made on any input; a punter may be granted free rides for the day at an amusement park since she’s the seven millionth visitor, for example. Finally, audit is the mechanism to log, or show, which actor has done what operation on what resource. Of these, audit is probably the most alien feature to the home user.

On the home network, a resource can also be an actor just as an actor can be a resource. On a business computer network, access control is often handled in a centralized fashion with a set of computers and software dedicated to access control. One popular access control protocol is Kerberos [NYHR05], which also the Windows Active Directory and Mac OSX access control subsystems are based on. A home network on the other hand may not have a central authentication server on site. Possible solutions would then be either to ignore access control altogether, let resources handle the access control themselves, employ a trust model, use an external authentication server as a service, or a mix of these.

User authentication on computer networks are usually done using one or more of the following means: knowledge-based, token-based and systems based on biometrics. Knowledge based authentication means that the user is able to recall something (i.e. a password) or recognize something (i.e. a picture). Knowledge-based and token-based authentication is often used in tandem, as in automatic teller machine (ATM) authentication. The predominant way of user authentication is knowledge-based: entering a password or a PIN.

The two fundamental types of access control are Mandatory Access Control (MAC) and Discretionary access control (DAC). MAC is an appropriate model for multilevel secure military applications, while DAC is used in the file systems of popular operating systems such as Linux. In Mandatory access control, ordinary users have no way to influence on the access settings

---

<sup>5</sup>Anderson uses *subject* to denote a physical person and *principal* to denote an entity – a person, a role, a piece of equipment – that participates in a security system [And01].

for any data they create; the access rights are handled centrally and are thus, from the user's perspective *non-bypassable*, *always-invoked* and *tamper-proof*. Contrarily, with Discretionary access control, the owner of a resource decides on its access settings. In the case of computers, operating system files are usually owned by 'root' or 'the system' and given minimal rights, if any, to other users. DAC rights are usually expressed through Access Control Lists (ACL), per-resource lists of actors and groups and their respective access rights to that resource. Due to their distributed nature, ACLs do not include a mechanism to categorize and group information objects [ZS96] throughout the home network, and in a large environment, the handling of masses of ACLs can be cumbersome.

A further access control type is Capabilities, where each program that will access any other resource is given the rights to do so by a set of 'capabilities'. This is upside-down from most other models where the access rights is a property of the data or resource *being accessed*. A real-world example of capabilities would be that anybody with a set of car keys would have the capability to use the car they fit.

An important concept on systems security is the Principle of least privilege. It states that no user, no program, no system, even no component (in short, no actor) should be granted more privileges than it needs for operation. Properly deployed and implemented, the principle of least privilege would do good for the security on the home network, when for example a compromised low-security temperature sensor in the sauna cannot be made to turn the sauna stove on or off. A corollary for the user is that the principle of least privilege can be trusted: a system cannot apply the principle at one place and not do so at another, and a system cannot *act* like it's abiding the rule when in fact it isn't.

A newer approach to expressing rights is Role-Based Access Control (RBAC) [FK92]. Users are allowed access to resources based on their roles within their organization, or seen the other way, a role specifies a set of transactions that a set of users can perform. A user may have multiple roles but depending on her work context, and she may need to select a certain role (or set of roles) for a specific context. The roles are managed centrally, so RBAC is in fact a form of Mandatory access control. Roles can inherit permissions from other roles (as defined in the  $RBAC_1$  model [SCFY96]<sup>6</sup>), so within a home network context, users of the 'parents' and 'children' roles can also have the

---

<sup>6</sup>The same text also defines the 'advanced RBAC models'  $RBAC_2$  which adds the *constraints* that one role may be incompatible with another – e.g. disallowing the same individual to be both the judge and the jury, c.f. the Chinese Wall security policy [BN89] – and  $RBAC_3$  which is the consolidated model of  $RBAC_1$  and  $RBAC_2$ . While it's debatable whether the advanced RBAC models actually are useful on the home network, they might be relevant in for object and environment roles defined in Generalized role-based access control (§ 5.5).

role ‘inhabitants’. Inhabitants would have access to enter the house at any time but only parents would have access to the bar cabinet, for instance. The access of roles should be built on the principle of least access, which supports the integrity of the resources.

Ultimately, access control is an issue between the devices on the network; a node communicating on the network may do so representing its user as well as doing automated background work. The authenticity of network communication, both in terms of its origin, its integrity and whether the request should be authorized, is a problem that needs to be tackled, though this hardly is a user centred security issue. The origin and integrity of a message can be verified either by direct point-to-point wiring or by using cryptographic check values, i.e. public key *digital signatures* or message authentication codes (MAC), used as a secret in symmetric encryption. Cryptographic measures, especially public key methods, may admittedly be an overly complex task for small nodes like light switches. Such devices could instead use IFF (“identify friend or foe”) type two-way authentication, where the device only answers if provided by a correct secret, hash-based access control which allows a master node to lock and unlock the node or lightweight encryption schemes such as NTRU or TEA [WSRE03]. Of course, in a wireless environment, unencrypted secrets could easily be eavesdropped.

Commands must be authorized before they can be executed, to safeguard for example from a rogue incoming message to the alarm system to turn itself off [Ell02]. Furthermore, the integrity of messages must be ensured and messages need to be protected from replay attacks. This can be done using digital signatures or message authentication codes in conjunction with message timestamping or sequence numbering, so it is more a problem of implementation than a technical one. Solutions like this still needs to take care of key distribution and be possible to build on ‘peanut devices’, expected to be common on a future home network.

#### 4.4.1 Motivation

Why employ access control at all? The most immediate motivation would be to keep confidential information on the network private, but this is just the start of it all. Access control should be applied to both internal threats as external ones, for a variety of reasons. One “threat” is that of the users themselves – even the most technically apt user sometimes makes mistakes [SG02].

In a networked home environment, there are large masses of data and media to protect, both in terms of privacy, integrity and availability.<sup>7</sup> There

---

<sup>7</sup>By data we mean the information the network needs to operate and by media the information that can be interpreted as music, photographs or video. In the words of

are users, devices, and processes that are, or are not, recognized and consequently should to an applicable degree be “trusted”, ie. authorized, to do or not to do certain things. While processes on the home network are digital and can be made identifiable, not all things in the home are (think restricted areas such as the medicine cabinet), and nor are the users.

Users’ data and media may indeed need to be kept safe, and in different ways to different audiences. Information about the whereabouts of the inhabitants of the dwelling should be kept away from unauthorized eyes. The family’s family photographs should be available for the family members and relatives; subsets of the photographic material may be available for friends or be completely public. Friends and relatives may be allowed to add meta-data (comments, tags) to the pictures while other may be allowed only to see the pictures, and so on. Some devices on the home network may also be shared to a wider audience. Intellectual property rights (IPR) issues aside, a family member might want to allow others to listen to the music she listens to, but not allow others to change the music. Both from inside and outside the home, public data should be kept from accidental or deliberate modification.

All this raises a relevant question: How much security do we need? Will things that are considered “secure” (whatever that means) in the current non-networked world be insecure once we move into pervasive computing? [Sta02b] Stajano presents the cynically realistic ‘Big Stick principle’ to help reflection on the topic: Whoever has physical access to the device is allowed to take it over [Sta02a]. The Big Stick principle is already applicable to many real-life devices such as the remote control at home or the refrigerator in the coffee room at the lab. Or seen from another perspective, if somebody has physical access to your laptop, she will be able to get into it<sup>8</sup>. Stajano suggests that in many real-life applications, the most appropriate protection seems to be social and territorial; even though a person, say a guest, *may* go and empty the fridge, bar cabinet or CD collection, social convention dictates that this is not appropriate behaviour and a guest doing so will not be invited to the home again. Clearly, the access control at homes need not primarily be targeted at people we trust (but we might want to apply logging to our system anyway, just as a measure of perfectly ordinary paranoia).

To understand why we need security, we need to understand what threats there are on the home network. Devices on the home network can be misused for a variety of reasons [SV04]. They can be used to penetrate the privacy

---

Robert X. Cringely, “Data is generated, media is distributed”

<sup>8</sup>Stajano notes that there may be exceptions to this example; the computer may have a password set in BIOS which is required for booting the machine, or the laptop’s hard disk may be encrypted so that it isn’t readable even if moved to another computer for analysis

of homes (generally not a very appealing reason unless the home belongs to a celebrity or the crush of an adolescent – still even people who aren't celebrities or objects of unarticulated affection should be allowed enjoy their rightful privacy). Devices can be broken in to for gaining more permissions through one device and then propagating the damage with elevated privileges throughout the network. Compromised home network devices can be used as attack vehicles for malicious activities such as denial of service attacks on other devices or other networks, spamming or being the host of morally questionable web sites. Apart from being an embarrassment to the home network's owner, such activity may harm the intended data on the device and it consumes the resources of the network. From the Internet service provider's point of view, a misbehaving home network is disruptive to their service and as such is eligible for being cut off from the Internet.

Traditional threats of computing such as viruses, worms and trojans still apply, and in a ubicomp environment, such threats could have devastating effects. Transactions could be eavesdropped or susceptible to man in the middle attacks. Teleworking may create further security requirements, especially if there are many teleworkers working for different companies in the home. Both in these cases and as we shall see in § 5.2, if such a device gets compromised, it may act as a bridge between one network and the other.

A device can be stolen, which has several implications. In increasing level of severity: the device goes, the data goes, and the access rights attached to the device goes. As the device may contain cryptographic material (keys), the loss of a device may result in compromise of the whole network.

A security threat salient to ubiquitous networks is location privacy. A ubicomp network will need to know the location of its users at all times, but this information must not be allowed to leak, even to the other users of the home if so needed. A thief would also be very pleased to know for a fact that there is nobody in the house. And it would be irresponsible for the architects of pervasive computing to build a world that could easily be misused as a surveillance infrastructure [Sta02b].

In a home setting, a relevant application of access control is that of convenience. Domestic appliances can be controlled [JSL<sup>+</sup>04, HT04] and adapt to their users' preferences and offer personalized services, such as a personalized view to a media library. A media player could play or pause the music depending on who's around. Aggregated with other contextual information such as the time of day or the user's schedule, the home network can help to realize the visions of smart houses.

### 4.4.2 Means and methods

There are different means for access control depending on what resources are to be secured and from what they are to be secured. The most common access control scenario to the user is that of logging on to a computer system with a username and a password. On the network, access control can be handled by a firewall and network address translation (NAT) to keep unwanted traffic out. A house practises access control on its inhabitant by using an out door equipped with a lock and key mechanism.

A user is usually authenticated using one or more *factors*. Practically any factor can be used but the most common factors are something the user *knows*, something the user *has* or something the user *is*. Of these three, the first one – i.e. a password or PIN the user knows – is the least secure and the least usable. A password can be forgotten, written down and accidentally or deliberately revealed to another. Something the user has, e.g. an ATM card, is the second worst in case of HCIsec; such things can be forgotten, lost, stolen or forged. Given the choice of one factor, what the user is – i.e., using biometric characteristics – would generally be the preferable option from a HCIsec standpoint [CAJ03].

Other, less commonly used authentication factors are the user's location (whether an “absolute” location as at a given place within a building, or relative to other users [MGH06]) and context (see § 5.3.1), time of day, size of transaction, whether the transaction was pre-authorized and cybermetric factors, i.e. using the computer's hardware and/or software setup (used e.g. by Windows Genuine Advantage [Mic06]).

Authentication based on knowledge can be divided into *recall* and *recognition*; the use of alphanumeric passwords and PINs are based on *pure recall* – assuming the user hasn't written the password down somewhere – while graphical passwords are either *recognized* or pointed on by *cued recall* (more on this in § 5.3). An item that the user carries (i.e. has) – an authentication token, a printout of one time passwords, a cell phone or an injected RFID tag [War05, Bah02] – can be used for authentication, provided the item can be kept strictly with its rightful user. Two important corollaries to this requirement are: the loss of a token must not result in unlimited misuse if the credentials stored on it [SV04] (i.e. revocation of rights must be efficient), and if a token is misplaced, the whole chain of trust can be rebuilt on the revoked right. Biometric information, measurable human physiological or behavioural characteristics such as a camera picture, a fingerprint or a voice sample can also be used to identify a user.<sup>9</sup>

The term ‘biometrics’ is used to refer to any and all of a variety of identification techniques which are based on some measurable physical, physiological

---

<sup>9</sup>It can be argued whether these characteristics are something the user *has* or *is*

or behavioural – and difficult-to-alienate – characteristic. Biometrics are used to identify and authenticate humans and is, within certain applications, both a promising and convenient way to do so. Biometric (“one-to-one”) *identity verification* works by comparing the user’s captured biometric data with the recorded ones. A current example would be the Finnish “biometric passports” which contain biometric data of the owner’s passport photograph. Biometrics can be classified as follows: [Cla94]

**appearance** the familiar passport descriptions of height, weight, colour of skin, hair and eyes, visible physical markings; gender; race; facial hair, wearing of glasses; supported by photographs;

**social behaviour** habituated body-signals; general voice characteristics; style of speech; visible handicaps; supported by video-film;

**bio-dynamics** the manner in which one’s signature is written; statistically analyzed voice characteristics; keystroke dynamics, particularly in relation to login-id and password;

**natural physiography** skull measurements; teeth and skeletal injuries; thumbprint, fingerprint sets and handprints; retinal scans; earlobe capillary patterns; hand geometry; DNA-patterns; and

**imposed physical characteristics** dog-tags, collars, bracelets and anklets; brands and bar-codes; embedded micro-chips and transponders.

An ideal biometric should possess the following characteristics: [Cla94, CAJ03]

- Universal – everyone should possess the characteristic,
- Unique and exclusive – each individual should have a unique version of the characteristic,
- Permanent through life,
- Indispensable – the identifier should be available at all times,
- Collectable, and digitally storable,
- Precise,
- Easy, efficient and not too costly to record,
- Convenient and fast to measure,
- Acceptable to contemporary social standards.

There is no basis for identification that fulfils all the abovementioned characteristics.

Biometric identification comes with a host of other problems, many due to the inherent uncertainty of biometrics: no two biometric samples from the same subjects will look exactly alike (within-subject variability), but two samples from different subject may be similar enough (between-subject variability) to confuse the system. Biometric information is not a secret. People leave fingerprints everywhere and there are irises anywhere one looks [JŘ03]. The biometric signature also needs to be tested for “aliveness” to verify it really is an iris or a thumbprint, not an image of one.

Biometrics are a promising way to conveniently authenticate users, but they cannot be used to authenticate computers or messages. Since biometrics aren’t secret, they cannot be used to sign or encrypt messages.

#### 4.4.3 Multi-factor authentication

Authentication can be made more secure or more certain using multiple authentication inputs. In traditional multi-factor authentication, several different factors are combined for increased security. An everyday example is payment with plastic; the user both *has* the right ATM card and *knows* the correct PIN code (this everyday example was an exhibition of *two-factor authentication*).

Multiple inputs can also be used to increase the certainty or *confidence value* of the authentication [CAMN<sup>+</sup>02]. This is especially true in cases of non-intrusive (e.g. biometric or RFID-based) authentication where one input may only give limited certainty of the user. A user could, with reasonable confidence, be deduced using the ‘constellation’ of things he is wearing or carrying, especially if those things were RFID-tagged. A similar solution could also be used in a distributed scenario where, for example, a house already knows that the users within are authorized, so a ‘peanut’ sensor only needs to sense that there *are* people in the room: by induction, they are authorized users (such a system could of course be foiled if the more secure authentication mechanism was bypassed somehow).

### 4.5 Making security usable

Though it was stated already in 1975 by Saltzer and Schroeder that usability is an important factor to computer security [SS75], usability and security are rather new acquaintances with each other. The hindering misconception has long been that it is not possible to make things easier to use while while simultaneously making them more secure to use, when the point really is

that we need find ways to make products secure *and* usable at the same time. Making an existing usable product secure or making a secure product usable is a complicated and potentially gargantuan task (often asked of designers to perform on a nearly finished product) that may end up a patchwork of ill-fitting components that is neither very secure nor very usable.<sup>10</sup> Designers must understand that security and usability elements “can’t be sprinkled on like magic pixie dust” – it is much better to consider the security and usability as aspects of a common goal: fulfilling the user’s expectations [Yee04]. Security just restricts access to operations that have undesired results, while usability improves access to operations that have desirable results.

Alma Whitten and Doug Tygar identified the following properties of the usability problem for security: [WT98]

**The Barn Door property:** If a secret leaks, there is little point in trying to secure it anymore as there is no way to make sure an attacker hasn’t intercepted it already,

**The Weakest Link property:** the security of a networked computer is only as strong as its weakest component<sup>11</sup>,

**The Unmotivated User property:** users generally do not use their computers because they want to manage security but rather send email, browse the web and so forth,

**The Abstraction property:** computer security management and policies may be alien and unintuitive for the general user, and

**The Lack of feedback property:** providing meaningful feedback to the user so she can check whether the set configuration really is what she intended.

Dourish et al performed a user study on users’ perceptions of security and the results were grim [DGdlFJ04]. Users see security as a source of frustration, they – especially younger users – are pragmatic about their security needs (meaning they do what it takes to get their job done and have no problem using a computer chock full of viruses if that doesn’t matter to their goals), and have an overwhelming sense of futility towards security, since “the unknown others (hackers, stalkers, etc.) will always be one step ahead”. Users perceive security as a barrier in both a user-hindering sense and as an all-encompassing barrier between them and the bad world out there. Thus,

---

<sup>10</sup>Admittedly, in the context of a large, already existing product, this exercise may be unavoidable. The task ahead is still potentially gargantuan; project managers take heed.

<sup>11</sup>The weakest link property can in the context of this thesis be extended to consider the whole network, not only one networked computer.

in the eyes of the user, spam protection, virus protection and protection of evil network activity is all security and is handled by a firewall. Security is generically something to keep things out, like a locked door.

While all-encompassing and general design techniques, rules and heuristics for designing things with usable security do not exist yet, there is ongoing work in the HCIsec research community to address these issues. Some guidelines follow: An application should come with secure defaults (for a secure “out of the box experience”). The security decisions the user has to make must be woven into the flow of the process (as opposed to interrupting it with seemingly unrelated security questions) and the application should guide the user into making secure choices. The most natural way to do a task should also be the safest. The user should be able to easily review and revoke any authoratory decisions [Yee04]. The application should tell the user what it’s doing if what it is doing has security implications [Hol99]. Though rather supporting than orthogonal to these guidelines, Alma Whitten proposed the following design principles and techniques [Whi04]: Security decisions should be presented *well-in-advance*, as opposed to the general user interface design practises where information and choices are presented just-in-time; users should be allowed to make security choices in *stages* (‘safe staging’) corresponding to his or her skills and acquaintance with the system, and by systematic way of *tailoring metaphors* (icons and so).

As far as methodology goes, both security and usability design advocate an iterative (not linear) design process. One methodology to create usable and secure software, AEGIS, is a well-defined blend of Contextual design and risk analysis with both facilitators, stakeholders and security experts as part of the development team [FSH03]. Contextual inquiry (and design) can be given a security flair by making sure the user interviewed does things that involve security. Regardless of interview technique, users cannot in general be asked about security directly; partly because users may be influenced to change their perception of security, and partly because we are trying to make security a seamless part of the application [SG02].

The Prime Directive, so to speak, of usable security is to make something both secure and usable at the same time. The motivation is simple: Security that is easy to use will be used, and security that is hindering won’t. Usable security should not get in the way of the user. Rather, it should try to assist and encourage the user into a more secure user experience. A product, be it a software product or anything else that has security implications, should guide the user into making secure choices. Conversely, it should be *hard* to make an insecure choice and the system has to communicate to the user that he is doing something that may be irrevocably devastating. Of cocourse, HCIsec takes its target audience into account; a user interface for the typical home user would be very different from that of an experienced

system administrator who plays with irrevocably devastating things for a living.

In fact, many of these guidelines are familiar from Nielsen’s usability heuristics [Nie93] discussed in § 4.1. They are just presented with a security perspective.

Security and usability may also need to be balanced to find a good solution: Even though a really secure but complex solution is traded for a lesser secure but unobtrusive one does lower the *theoretical* security of the system, the *effective* security is increased just because it is used.

Security should be *implicit*, i.e. be built into the tasks the users have [SG02]. For instance, if a user wants to share a file, all the steps of doing so must be included in the actions the user needs to take when sharing that file. There are two prototypical paths a user may take to share a file. Either she marks the file or folder in which the file resides readable by target user (in case the target user is not a member of the system’s security domain, “everybody”). She must make sure that other files aren’t inadvertently exposed, and she must remember to remove the sharing after the target user has retrieved the file. The other way is to send the file by mail (which most users *perceives* as secure) or, as is getting increasingly popular, using an instant messaging program (which at least may be encrypted). Obviously, the second way is much more clear to the user: select a recipient and fire away. It’s a one-shot process and there’s nothing to clean up afterwards. System administrators in business environments on the other hand hate this approach: sending an identical file to multiple users by mail wastes resources and if all the recipients make changes to the file, it is impossible to see which one is the most current one. Clearly there is space for a way for the user to accomplish her goals with a transaction that is both easy, safe and electro-ecologically sound.

To take the idea of security not getting in the way of the user to the extreme (and spice it a little with the idea that a bit of security that is used is better than a lot of security that isn’t), there is a lot of security that can be automated. Devices can be set to auto-configure themselves securely on a home network [SJF<sup>+</sup>03] or an email cryptography can be configured to retrieve unknown public keys from a server [SG02]. A laptop computer can use an encrypted file system without the user even needing to be aware. Even though the result may have loopholes, it is better than no security, as long as the user isn’t lulled into a false sense of security – which in itself is a lot more dangerous than knowing you have no security.

## Chapter 5

# Previous work

In this chapter, we present research, findings and information on security and access control which can be adapted for home networks. It was previously noted that not many solutions exist for access control specifically designed for domestic use. Security and access control, even user friendly access control, are well published topics in academic circles in any other field than the home. Whether this is because the home is such a complex environment or because researchers spend way too little time there shall be left as a subject for future research. However, findings and research from related fields, most notably that of pervasive computing, should be applicable on the home network environment. A lot of previous work has also been presented elsewhere in this thesis.

### 5.1 Security on home computer networks

Users on home computer networks have a fair share of guides and checklists to turn to in helping them create and maintain a relatively safe environment. The Computer Emergency Response Team (CERT) maintains a document which acts as a primer on home computer network security, covering both the risks and the countermeasures associated with Internet connectivity [CER01]. For the Finnish audience, a large number of state and private organizations have teamed together to support the sites [Tietoturvaopas.fi](http://Tietoturvaopas.fi) and [Tietoturvakoulu.fi](http://Tietoturvakoulu.fi) (*Information security guide* and *Information security school*, respectively). Tietoturvaopas publishes information for home users about computer security, cell phone security, spam security and privacy. Their most relevant publication for home users is the eight page leaflet *Joka kodin tietoturvaopas* (“Each home’s information security guide”) [Tie]. Tietoturvakoulu is aimed at children of school age and their teachers and has paths for both younger and older pupils.

The computer security guides usually give the following advice: Keep your computers' software up to date; patches can be found at the software manufacturer, use automatic updating if possible. Use a firewall. Use virus protection and keep the virus detection profiles up to date. Beware of suspicious-looking email and web sites. Do not reply to spam and don't give your email address to suspicious web sites. Set email reader and web browser security settings to maximum. Don't run programs of untrusted or unknown origin. Make regular backups of your data. If you have a wireless network, use encryption to keep uninvited guests out.

The Finnish Computer Emergency Response Team (CERT-FI) publishes an advisory bulletin of vulnerabilities and a security blog<sup>1</sup>, and the Finnish Ministry of Finance publishes security guides<sup>2</sup> (The user's information security guide, Protecting against malware, etc). These resources are primarily aimed at the technical user.

## 5.2 Identified problems on pervasive networks

There are many challenges intrinsic to home networks, which are of pervasive nature. These challenges are foremost of technical nature, but naturally have implications on the user experience.

Architects of products and solutions that are intended for the home network often make false or oversimplistic assumptions about home networks. For instance, security products for homes are often designed like the home has only one inhabitant [E1102]. Sometimes, solutions which work well in an office environment are just copied to the home environment without taking into account that they are going to be used in an environment quite different to the office.

In a home network environment, devices can be connected to different parent networks through different gateways: mobile phones to the 3G network, some wireless devices to a commercial ubiquitous Wimax network and others to the domestic network. The home user can have his wireless device connected to a foreign network while accessing both that network's and his home network's resources. A service person visiting the home can be connected with his wireless device through his employer's cell phone connection while simultaneously being connected to the home network due to his service duties [SV04]. Cases like this create challenges for perimeter security as the perimeter no longer can be simply defined as the outside and the inside. The "outside" user is already on the "inside", and the traditional concept of a

---

<sup>1</sup><http://www.ficora.fi/suomi/tietoturva/varoitukset.htm> and <http://www.ficora.fi/suomi/tietoturva/cert.htm>

<sup>2</sup><http://www.wm.fi/vahti>

firewall becomes fairly useless.

Users can have no guarantees of the security of a foreign network. Privacy, including location privacy, becomes a problem when mobile devices poll their surroundings on a foreign network. This capability can also be exploited by intruders or malicious insiders. Users may not even be aware of the collection of their personal data. Villains could also exploit networks by injecting misleading information, stealing or tampering with electronic assets or disrupting critical services. The environment would have to take care of its access control needs on a much more distributed basis. It's a complicated mix of a common policy and "each device to its own".

To allow for such settings, all devices would need to agree on a compatible security policy, common trust management and a common protocol. There also needs to be a flexible and convenient method for defining and managing security policies in a dynamic and flexible fashion to support the users into making and maintaining good security. The home network also needs to support single sign-on and be robust enough to have its management functions always and globally available [HSU04, SGTGI04, CAMN<sup>+</sup>02, SV04].

The home network itself needs to be safe from physical harm, both in form of intrusion by malicious individuals and from the forces of nature. If it rains, close the windows. If the sun gets too hot, close the blinds (which is a more economical option than to turn the air conditioner to a higher setting). Factor in contextuality: if the house is alone, there is nobody to care whether all blinds are shut and it's near pitch black inside, as long as it makes sense energy-wise. And all principals – be they humans, machines that represent humans or machines don't – that are not (yet) mutually trusted need to be guarded from each other.

There may be many small, wireless, battery-driven nodes on the network, low on computational power (Frank Stajano calls these 'peanut devices'). Especially this kind of devices need to balance the conflicting demands of actively surveying their surroundings while using their power conservatively. Such nodes are also susceptible to DoS or *sleep deprivation torture* attacks from malicious or badly operating nodes which would drain the battery of the 'peanut node'. Like all radio devices, a wireless device's radio communication can be jammed by another transmitter on the same frequency area [HSU04, SA99, Sta00].

Finally, a problem most home users (and some corporate users) don't think of until after disaster has struck is backup of data.<sup>3</sup> The issue is further complicated by the fact that the users' 'data space', while it may appear uniform, may actually be distributed among a lot of machines both in and outside the home. Outsourced file storage can be backed up by the service

---

<sup>3</sup>This includes yours truly, in a near-final stage of this thesis.

provider, but backups at home must be automated. Also the interface to restore data must be designed for home use. One interesting solution suggested by [SV04] is that the backups are redundantly distributed among different devices on the home network, so that the data can be retrieved even if some devices carrying the backups should be unavailable.

### 5.3 User identification and authentication

A problem with the traditional means of authentication is that it doesn't scale well. In a work environment, it is only a minor inconvenience to log in to a computer by username and password but in a home environment of several hundred nodes, the old solution is no longer usable. And even where passwords can be used, there may be problems. Passwords are not easy to enter on devices with modest input capabilities such as cell phones or remote controls, or by children or elders with reduced motoric or cognitive skills. From a security standpoint, any password that is easy enough to actually remember is weak enough to be broken within minutes or even seconds.<sup>4</sup> 'Logging on' by traditional means is not in line with the activities on a home network. It is disruptive and it is personal, not a shared activity like many activities in the home [Bar05].

Instead, we need to look for alternative solutions, more suitable for the home environment. Some of these solutions can be found in the pervasive computing research. The following is not.

**Authentication by images:** Instead of authenticating with a password, the user could authenticate herself by identifying pictures [Dha00, DP00] or pointing on specific locations of pictures, coined 'passpoints' by their creators [WWB<sup>+</sup>05]. Using graphical passwords is based on recognition or cued recall, which humans are rather good at, rather than 'pure recall'.

To create a picture "password", the user selects a given number ( $P$ ) of pictures, either photographic pictures or algorithmically created 'Random Art', into her 'image bank'. To authenticate, the system then presents a set of pictures and the user has to pick which of the ( $T$ ) pictures belong to her image bank. When  $P = 5$  and  $T = 20$ , there are already more combinations than with a four digit PIN. In a user study, Dhamija and Perrig found that users were remarkably good at authenticating with pictures, having created them in a test one week earlier [Dha00]. Creating a graphical 'passpoints' password is a similar procedure: the user either selects or is selected a

---

<sup>4</sup>As a countermeasure, some systems freeze the account after a given number of unsuccessful login attempts. This only creates another problem, namely that of a denial of service (DoS) attack – any malicious entity could lock out *another one's* account.

picture. This picture needs to have enough characteristic details. The user then clicks or taps a given number of points on the picture. The user then trains herself in tapping on the right passpoints in the correct order until the system considers the recognizing to be stable enough.

**Access by proximity** Different solutions have been envisioned and some even tested for identification and authentication by proximity. Corner and Noble suggested an authentication token with which the user authenticates, wirelessly and with “zero interaction”, to a laptop computer, which would lock when the user disappears from the vicinity of the laptop [CN02]. Nakajima and Satoh take the vision further by having the user carry personal information and preferences in a ‘personal home server’ [NS04]. The name implies that the device is for personal and domestic use, and that *it* serves the home with its carrier’s information and preferences.

There are two immediate shortcomings in both of these ideas. First, the user must wear an identification token at all times if she wants her domotics to behave in a personalized manner. Secondly, if there are two identification fobs in the vicinity of a single-user appliance, which one should the appliance take into account? Hannu Kari has a suggestion which would solve the second one of these problems; using an identification token which is in galvanic contact with the user (such as a watch), the user would first identify to a device by touch – the skin would act as a transmission medium – after which the token would communicate wirelessly with the device as with the “zero interaction” token above [Kar06].

The first problem, that of carrying an identification token at all times, would be solved by injecting the identification token into the user’s body [War05, Bah02] but the author finds it highly unlikely that such a solution will gain high acceptance with the paying populace.

### 5.3.1 Context based authentication

Several authentication methods have been suggested that include the user’s context as input. Noda et al included the user’s role, location (‘presence’) and calendar information for arbitrating the access decision in an RBAC policy engine, using RFID technology for the users and a hidden Markov model to assist with incomplete sensor signals [NTH<sup>+</sup>06]. Incorporating trust in the model, the system could support users based on confirmed attributes such as the fact that they are employed by a business partner or public authority without having to list each of these users separately. The team also discussed that in a case (“context”) of emergency, the access control system would for example allow giving out personal and medical information when in an emergency room and any responsible doctor is in the same room.

A related use of the term context (based) authentication is using shared context between devices to create shared secrets. These shared secrets can consequently be used as cryptographic tokens for creating secure channels [May06]. One application of this could be this: when the microphones of several laptops “hear” the same signal, they can be assumed to be in the same area.

## 5.4 Device management

Only authorized devices should be allowed to join in a user’s home network. This implies that only authorized service providers should be allowed to register services on the network. On one hand, a user would not have to have a malicious device on their network, but also a device belonging to a neighbour’s network should not show up on the user’s home net. Also the de-registration of devices should only be allowed from devices authorized to do so. Whether we will allow a device on our network, and whether processes on it will be authorized to interact with the network, depends among other things on whether we ‘know’ the device and where the device is [SV04].

One method for device authorization in a rather unobtrusive fashion is suggested by Stajano and Anderson in [SA99] as The Resurrecting Duckling security model. Just as a duckling is imprinted with the first mother duck candidate it sees, a device gets imprinted by a master device whereby it is said to get a ‘soul’. The imprinting, akin to Bluetooth device pairing, is either initiated by galvanic touch of conducting elements (i.e. electric contact) or by a short range wireless link that cannot easily be intercepted by an unwanted third party. After being imprinted, the ‘duckling’ device will only trust its ‘mother’. The duckling may still interact with other devices, it just cannot be *controlled* by them. De-registration of the ‘duckling’ device can be configured so that it only can be requested by the ‘mother’ device, which leads us to the “resurrecting” part of the security model: as the duckling “dies”, its soul dissolves and its body returns to its pre-born state, ready for another imprinting that will start a new life with a new soul. Of course, the device can also be configured to de-register by any identifiable transaction, or by a simple timeout (so that the duckling dies of old age).

This secure *transient*<sup>5</sup> *association* has some very desirable functions in a domestic setting. A cell phone, a laptop or an iPod may be imprinted to its owner and refuse to operate if stolen, i.e. with a foreign ‘mother duck’. On the other hand, a home’s back door will open with a universal remote control that is imprinted with the house, but not with an identical one brought by a burglar. And because the association is non-permanent, a device (even

---

<sup>5</sup>fleeting, temporary, non-permanent

the universal remote control) can be sold or given away when it has been properly dissociated.

This model works fine in a centralized, master-slave type environment, but in a home environment the model can get restrictive. If a duckling also were allowed to be a mother duck (“have offspring”), then the hub-like orientation would change into a hierarchical “family tree” where ducklings would heed not only their mothers but also – and with greater respect – their grandmothers, and ancestors before. This policy was suggested, though in lesser poetic wording, by Hannu H Kari [Kar06]<sup>6</sup>, while unbeknownst to him, Frank Stajano had returned with his resurrected duckling model in a way much more suitable for ad-hoc networks [Sta00]. By changing the duckling’s requirement “but it can’t be controlled by them” to “. . . it is happy to talk others, and even obey their requests, as long as mummy said it was OK to do so” [sic], approved ‘ducks’ (“relatives” perhaps) may be allowed to upload new instructions to the duckling. Such instructions may be divided into low integrity and high integrity commands, so that a malicious “godfather duck” would not be able to ask the duckling to kill itself and then take over the duckling by imprinting. Conversely, the mother duck could backup her own soul in case she gets hurt, which otherwise would render a duckling uninstrutable and unimprintable (relevant, for example, if the mother duck is a remote control to the A/V system and the owner of it has a dog or toddler that likes to chew on remote controls).

As a practical implication, now one remote control (mother duck) can control all of the whole A/V system of ducklings, but the A/V components may also interchange information like aspect ratio of the picture or surround settings on the audio processor. Foreign ducklings could also be serviced, so that a friend’s camera could get location info from his friend’s GPS-equipped car. The flexibility induces new problems: the duckling should be able to shield itself from a sleep deprivation torture attack (see § 5.2) and other DoS-related activity from malicious entities, and from the other side of the transaction: should the camera trust that the GPS location is accurate? In case of a friend-to-friend scenario, a grandmother duck could have instructed her progeny that friends can be trusted.

**Secure discovery** To hinder an attack on services of the home network, discovery of services should be secured. If services do not answer to unauthorized service discovery requests they are harder to target [CGR04]. Services could also selectively disclose their offerings based on what credentials the user (or device) requesting the discovery has. Service discovery responses

---

<sup>6</sup>HHK likened the process of one device recursively authorizing the next one by touch to how an infection spreads. One is tempted to coin this method ‘The infected duckling’, weren’t it for the rather grim connotation.

that are sent unencrypted can of course also be eavesdropped. While secure discovery is not a part of the UPnP specification, it is easily added as an implementation of the protocol. Encryption is not used in UPnP service requests, but it is in event notifications so this could be a future addition to UPnP. Also, since UPnP builds on XML, it is perfectly legal to add this feature to the already existing protocol.

A variation on selective disclosure is that the identity of the service is hidden. It might be possible to discover a service but not the provider that offers it. In a distributed network, such a feature would actually be a benefit: with many services, e.g. a clock service, it is not relevant who provides it as long as somebody does.

Finally, a rogue device may announce that it provides services, which induces yet another problem in the equation. Fortunately, these can be handled by the backline technology and need not bother the user.

## 5.5 Role-Based Access Control and GRBAC

Traditional Role-based access control (RBAC) [FK92] revolves around *roles* that *subjects* (actors) have.<sup>7</sup> Each subject has an *authorized role set* which contains all the roles the subject can *enter*, i.e. all the roles the subject *possesses*. Furthermore, an RBAC system has *objects* (resources) which can be *transacted* upon. All transaction permissions are associated with roles, not subjects, implying that each role has an *authorized transaction set*. So, to execute a transaction, a subject must demonstrate possession of some role which is authorized to perform that transaction. As an example, a ‘person’ (subject) can be authorized to ‘unlock the out door’ (transaction) of the ‘home’ (object) if she is an ‘inhabitant of that home’ (role). Roles can also be hierarchical (defined in RBAC model  $RBAC_1$  – see footnote in § 4.4); therefore a parent (sub-role of inhabitant) may also unlock the door of the home.

Generalized role-based access (GRBAC) is an extension of RBAC [MA01]. It adds some contextuality to RBAC, giving roles not only to subjects (which in GRBAC are called ‘subject roles’) but also to objects (‘object roles’) and the environment itself (‘environment roles’). With GRBAC you can state a security policy to prevent the kids from turning on the television before they’ve done their homework (provided you have a way to actually test this), or after 23, or turning on the naughty channel.

Subject roles are essentially the same as RBAC ‘roles’. Object roles classify

---

<sup>7</sup>Despite its name, Role-based access control is only concerned with the authorization bit of access control, and omits the identification, authentication and audit aspects. It would thus be more correct to call RBAC Role-based *authorization* control.

the object. Pictures of the family can have object roles ‘family pictures’ – which itself can be a subclass of ‘photos’ – ‘2006’, ‘print this picture’ and ‘taken with in-house equipment’. As an implication, all objects need to be classified or the policy must specifically include the case ‘all objects’ (i.e. ‘not classified’). In some cases object roles can be computed, or provided by the provisioner of the resource. Finally, environment roles describe the state of the environment. ‘Weekend’, ‘network storm’ and ‘it’s raining’ are examples of environment roles. Since many environment roles can be active at the same time but not all environment roles are relevant for the access decision, the calculation and application of the right set of environment roles must be done carefully.

The GRBAC *transaction* is defined by the tuple  $\langle SRole, ORole, ERole, Op \rangle$ , i.e. an operation where a subject acting in role *SRole* performs an operation *Op* in the role *ORole* under the environmental conditions *ERole*. A *policy rule* is defined as  $\langle Transaction, PermissionBit \rangle$ , where the *PermissionBit* indicates whether the Transaction is allowed or denied, supporting both positive permission models (“everything is denied unless specifically permitted”), negative permission models (“everything is allowed unless specifically denied”) and mixed permission models (highest precedence rule decides). A beverage will be rewarded the first person to tell the author they’ve read this sentence. In addition to the ‘all objects’ pseudo-role, the (un)specifiers ‘all subjects’ and ‘all operations’ can be used to patch for missing role information in transactions.<sup>8</sup>

GRBAC has been proposed for use in future domestic applications [CMA00]. Using GRBAC, access control should be possible to present in a way, and in a language, that is comprehensible to the non-technical user. GRBAC also easily bends into accepting ‘confidence values’, calculated values and sensor inputs, allowing a non-intrusive access control process.

## 5.6 Security mark-up

Herein are described two XML based security markup languages that could be used for a home network backline.

**XACML** Extensible Access Control Markup Language is a specification by the Organization for the Advancement of Structured Information Standards (OASIS). It defines an access control mechanism based on XML documents. The specification defines three basic documents: Policy (description of the permissions), Request (a request for some Subject over some Resource

---

<sup>8</sup>The GRBAC specification does not mention a ‘all environments’ pseudo-role, but through extension, such a feature could easily be added to the set.

to perform an Action), and a Response (the results of applying the policy to the Request document). This implementation allows loading policies and requests from its XML representation, and perform an evaluation applying the policy. [XAC03]. XACML has a profile for expressing policies that use Role-Based access control (RBAC, see § 4.4).

Since XACML is an XML, it could be used in conjunction with Universal Plug and Play (UPnP), itself an XML implementation.

**XRML** The Extensible Rights Markup Language is an XML to express digital rights. It is useful for content publishers and distributors to state a digital rights management (DRM) declaration to the related media, but also applicable for Web Services. XRML is designed to be applicable for both simple and complex rights expressions in any stage in a business, workflow or business model [XRM01]. Microsoft has adopted XRML for its DRM process.

DRM protects the media from its users, not the other way around: Digital rights management is not to protect the user or the network but to protect the media. Generally speaking, DRM is in the interest of the media publisher, not the media consumer (to which DRM may even be hindering). From the consumer side, DRM has often been seen as overkill and unfair. With the Internet becoming an ever more relevant distribution channel for “regular users”, users also become media publishers, so it is possible that they also would become interested in securing their media using DRM techniques. However, DRM is outside the scope of this thesis.

**CBAC** This chapter cannot be closed before making a note that a ‘Context-Based Access Control’ does exist, but despite its fitting name, it is not an access control means for context sensitive home network applications. CBAC is rather an RBAC for firewalls. Different kinds of traffic can be filtered or passed through depending on application layer protocol state information (unlike traditional firewalling which was limited to examining the network and transport layers) as well as whence and where to the requests were going. Modern ‘personal firewalls’ tend to be Context-based.

## Chapter 6

# Access control at home

One of the problems with creating solutions for home network is that designers do not have a proper understanding of the home network environment and its users [E1102]. Similarly, users of the home networks do not have an understanding of what solutions a home network may offer. To shed some light on the reality, the author participated in two series of interviews: one with actual end users of home networks as they currently exist, and another set with different experts in the field. A state-of-the-art method designed specifically for gathering information from users in their right context is Contextual Design [BH99], which was chosen as the basis for the research for this thesis. The inquiry, interpretation and modelling steps were used and adapted to suit the realities of the research.

### 6.1 User interviews

In order to better understand the future uses and users of a future home network, a research team of the InHoNets project [inh06] performed a series of field interviews in *contextual inquiry* style. Along with the lines of Contextual design, the interviews were conducted in the families' homes. The main beneficiary of the interviews and the processed results of these were the project itself. We wanted to get some insight on the following things:

- what things reminiscent of a future home network people have, who uses them, how and why,
- what the people's attitudes towards this technology is,
- what people's attitudes are towards security and privacy are in their current lives, and

- what people do with their leisure time and how home network technology could support these activities.

As part of the second point, how people perceive technology, we also wanted to know what they would like the technology do for them (and how), and in which way they feel technology is helping or hindering them. By this, we discreetly enquired their attitudes towards computer security and security in general, since we felt that security and privacy, and in particular access control, are concepts that may be alien to our everyday users if presented in such terms.

We wanted to interview “real” and “realistic” families with teenagers, since such families are supposed to be the most interesting, complex and ignored ones from a computer security perspective [Eli02]. Five families were interviewed, all from Espoo and Helsinki, all living in different suburb areas. Two of the families interviewed were single-parent families,<sup>1</sup> the rest being of the more traditional two-parent kind. Each family had a teenage girl; in some families there were sibling brothers and/or sisters and in one interview, a friend to the teenage daughter attended.

While the method we used is based on Contextual design [BH97], several discrepancies exist between “pure” contextual design and our approach. Since we were interviewing families and using their limited and valued spare time, we only planned to use one hour of time per interview. Usually, an interviewer spends at least a (work) day with the interviewee. To get a full and realistic, *experiential* view of all the family members’ home and leisure time activities, a lot more time would be needed and significantly more data would be gathered. This, however, would have been considerably out of scope for our research, both in terms of material (more than we needed) and resources (more than we would have), a good overview of the situation served our purpose.

The creators of Contextual design note that our implementing of their methodology is an approved behaviour; Beyer and Holtzblatt say that their method should be taken as a basis and always adapted to the context of the work [BH97].

Two or three researchers were present at each interview. In accordance with the Contextual design method, one researcher asked most of the questions while another took extensive and detailed notes, and filled in with appropriate questions in between. We also took a lot of photographs from the interview sites. First, we asked general questions about the family and their activities with as much of the family as possible. After this, we let the individual family members show parts of their “home network” that mattered

---

<sup>1</sup>This is to say, the parents were divorced and the child/ren lived primarily with the one of the parents.



Figure 6.1: Our affinity wall in mid-progress

most to them. In particular, we wanted to give space for the teenagers to express themselves. In Contextual design spirit, we asked the family members as far as possible to *show* us the technology they were using and how they were using it, rather than just talk about it. We had agreed to stay no more than one hour at each family, but at each site, discussions were so lively that none of families wanted us to leave after our allotted time was up.

After each interview, the interview material was disseminated by the researchers present in the interview. In these interpretation sessions, we went through the interview notes line by line and made a sticky note for each unit of information. The notes were then organized on an affinity wall with headings for each logical cluster. Each family was ‘interpreted’ with differently coloured notes, which made combining and consolidating information easier (Figure 6.1, two of the five family interviews have been analyzed).

There were three outputs of the work. First, a spreadsheet was created which contains all the findings from the interviews (i.e. the raw data organized). Second, a fictional family story was written which represented the typical family, their activities and their technology, based on the interview findings and thus, reality. Third, a use case document for home network requirements was created.

### 6.1.1 Findings

The level of familiarity with and existence of traditional personal computers and computer networking technology varied greatly between families. One family did not have a home computer at all while another family had more computers than family members, wired and wireless Ethernet (*“with all encryption there is”*) and a stand-alone firewall/router based on BSD providing a permanent VPN connection to work.

Typically, there would be at least two computers in the home, one “main” computer which all family members would use and one older computer the children shared. Regardless of the technical know-how level, the prevailing view on home technology was that it should be easy to use and not get into the way of the users’ actual goals. While one of our system administrator types clearly enjoyed having built this technology, maintaining the system was never one of those goals (only the inevitable means for the administrator to allow the family to use the system). It also generally wasn’t very important how something worked, as long as it did work. Thus, one family didn’t really think they had a filing system for their music and photography, “iTunes and iPhoto takes care of it all.” This completely supports the findings and views of literature.

Some of the most important applications of the technology were communication and entertainment. Communication was done over a multitude of channels: voice and text over cell phone, text and pictures over email and over instant messaging, which was heavily employed by all the teenagers using computers. The people communicated with their friends and relatives, hobby associates and in work related issues. Notably the teenagers interviewed which, as noted above (§ 6.1), all were girls, used a lot of time on communicating with their friends and very little on gaming. We hypothesize that the emphasis on communication in favour of computer gaming may be a gender issue among the teenagers (which is consistent with [Gro04]).

The most important entertainment manifestations were photographs and music. Generally, at least one digital camera and at least one portable digital music player existed in each home, in addition to cell phones capable of both rudimentary imaging and music playing. Photographs are taken and stored, and, interestingly, to a lesser degree, watched. An explanation

to this behaviour was given to us by one interviewee who is a researcher in sociology: people mark the importance of a situation by taking a photograph of it. Photos are also shared, usually by emailing but in some cases through a service on the web. Almost nobody had ever made hard copies of their digital photographs at a photo shop or printing service. Teens also took a lot of pictures of their friends, hobbies and other interests with their cell phone cameras, but nobody had managed to get their pictures off their phones and onto their computers (though all had tried).

Portable music players were present in all families; in most families more than one. Usually, each family member had their own music collection and to some degree, family members shared music with each other. Sharing music over “unapproved channels” was surprisingly rare. Instead, families copied their favourite CDs to their computer music libraries or bought music from online music service providers. Portable radio/CD “boxes” were also very common. We did not witness any digital video recorders (PVRs) for recording TV programming in any of the families, though there were a few VHS video tape recorders present, usually in low utilization. Contrastingly, one family used a VHS video tape recorder quite a lot to watch store-bought movies as a family activity. Others had DVD players or gaming consoles doubling as DVD players, also used to watch rented or store-bought movies. A few families had a gaming console but it was in active use in only one family. Elsewhere, teenagers typically liked to play small Flash-based games over the net.

All families with computers had an always-on broadband Internet connection. Nobody seemed very concerned about the technical specifics (such as bandwidth) of the connection, as long as it just works – which it usually did. Most families also had a printer and some had scanners. What we did not find in nearly any installation was sharing of resources: one family had a disk space shared on the network but nobody had made the printer available from all the computers. If somebody wanted something printed, the usual mode of operation was to copy the file onto a USB memory and print the file from there. The usual explanation for this mode of conduct was that it served them well enough so nobody had bothered in ‘fixing something that already worked’.

While some had access to their work data from home over the network, nobody had made data on their home network available from outside the home. The usual explanation was that it was complicated, insecure or both (or they felt that they wouldn’t be able to do so in a secure manner), or that the families simply hadn’t thought about the option. The ways to get data from a home computer to work was to send it by email to oneself. In fact, we saw evidence where email was used as a remote-accessible file store with the added bonus of having metadata such as time and context with the file. USB

memory drives were also used, but to a lesser degree. If somebody needed a file at home posted to work when they already were at work, he or she would call home and see if anybody could do the favour of mailing the file, or that failing, just blame oneself, not take it too seriously, and mail the file the next day. Related to this, we interviewed another family for another project that had their lighting system controlled by home automation. While the lights could be monitored and to some degree even controlled from outside the house, the parents of the home felt that no home automation more critical than lighting would ever be remote-controllable in their house, due to security risks.

One family member usually had the technical responsibility of the home network, the adult with the highest technical understanding. Regardless of the level of technical know-how of this person, the job was seen as cumbersome and something done when there *really* wasn't anything else to do at home. Security and computer usage policies did exist but they were usually unstated and implicit, and varied quite a bit from the highly security conscious to the nearly fatalistic "I know this computer is full of viruses but it doesn't really matter as long as we can do our stuff with it. When it gets too slow and infested, we just install it again." Another pragmatic approach was "We have Macs, so we haven't really got a virus problem." A remarkable exception to the point was wireless networks: only the family that had a highly knowledgeable administrator had a wireless network, which had "all security applied". The rest did not want a wireless network because it was perceived insecure.

We also enquired how the people regarded and protected their real-life privacy. The most immediate demand for privacy was that of physical mail<sup>2</sup>. The interviewees also strongly considered their homes private places, off-limits from prying eyes, though they did not regard others "spying" on them (by looking through windows) a real threat, with the explanation that "why should anybody be interested in us?"

Families handled access control in different ways. Some devices and data were seen as personal and would not be touched by anyone else except when explicitly given permission by the owner or with the understanding that if a parent suspected something bad, s/he would have the right to take a look, but even so, only with the child present. This was true both in case of digital data, physical mail or a personal (physical) photo album. In some families, family members had personal user accounts on the main computer. The explained rationale was more to keep documents and settings out of the other users' sessions than that of document privacy. Elsewhere, family members had named folders for their personal data which the other family members would stay away from by mutual agreement. One family liked to share

---

<sup>2</sup>There is a strong tradition for mail privacy in Finland

photographs with their friends and relatives abroad, especially as Christmas greetings. Such greetings were posted on publicly available services on the web with hard-to-guess identifiers for an outsider. Finally, we witnessed a very pragmatic approach to monitoring and regulating the children's network usage: only the computer(s) in public space had an Internet connection. A more technically inclined variation of this was that the children's Internet connections were regulated by time based rules in the firewall – no Internet after bedtime.

## 6.2 Expert interviews

The author conducted three open-end interviews with different experts of the trade: Ben Lavender is a technology manager at the BBC, responsible (among other things) of the iPlayer project, an Internet-equivalent of the video tape recorder; Risto Linturi shared his views on access control in domestic building technology and home automation, and Hannu H. Kari, mobility professor at the Helsinki university of technology, talked about security considerations in ubiquitous home networks. In the end, the third of these interviews proved the most fruitful for this thesis. Mr Lavender was interviewed over the phone, Mr Linturi over email and Mr Kari in a face-to-face interview.

**The BBC iPlayer project**, previously known as the iMP (interactive media player),<sup>3</sup> was a large scale test of a peer-to-peer structure to distribute selected TV and radio programming of the BBC. The BBC is hoping to release iPlayer for a general UK public in early 2007 – the biggest hurdles for this are now political, not technical. Due to TV licenses, the iPlayer for television programming would only be available for the UK viewers, though Lavender doesn't regard a worldwide service comparable to BBCworld impossible, given an applicable business model, for example commercials and a subscription model. iPlayer for audio content would probably be offered worldwide soon after the public iPlayer launch. The iMP ran two trials, trial one with 1000 early adopters, trial two with 30'000 users with a demographic corresponding to UK broadband users. Users would also be able to copy the programming to their mobile devices. The iPlayer also included a simple recommendation engine.

The only aspect of access control the BBC were concerned with was that of making sure only UK viewers, i.e. those who've paid the compulsory UK TV license, were able to view or hear the offered programming – TV piracy is taken very seriously by some, and is a hot political issue. The required form

---

<sup>3</sup><http://www.bbc.co.uk/imp>, probably to re-appear as <http://www.bbc.co.uk/iplayer>

of access control would mainly be executed with digital rights management (DRM) techniques and verifying that the user has an IP address in the UK.

**Domestic building automation** has a host of inherent problems. Risto Linturi traces the origins to most problems to the builders and contractors of the automated apartments; it is they who order and specify the systems deployed, not the people who actually live in these apartments. While building automation can help to reduce living costs through the control of heating, water use, air conditioning and electricity use, the builders are more interested in what they can get for less and the contractors in what they can profit in maintenance fees. Such technology is often proprietary and is hard or expensive to expand later on, and locks in the user with a specific vendor<sup>4</sup> Also, inhabitants do generally not yet have a comprehensive view on what home automation can do for them. All these factors combined result in slow adoption – and slow evolution – of house automation. Another issue in which house automation could be a tremendous help is in storing the plans, simulations and all service documentation of the home in a digital and searchable format that is automatically updated through logging and accessible to service personnel. This kind of service is typically available only in new commercial sites.

Different users of the network thus have and need different roles. In a scenario where a service company takes care of the maintenance of an apartment, a service person should be allowed to calibrate a sensor that the inhabitant should at most be allowed read access, due to responsibility issues. Some devices, for example a movement sensor, provides input to number of other subsystems such as heating and security (fire and burglary) but these subsystems may have different access groups and produce event data to different service providers. The user may want to add capabilities to the home network, such as sauna control or integrate the entertainment centre with the home automation, and of course the user should be allowed to do that. This calls for interoperability but also diffuses the question of responsibility.

An automated home is a complex engine with hundreds of nodes. It must be very reliable and there must be ways to report and repair when something goes wrong. The house automation system is also uplinked to the service provider(s), a connection that must afford confidentiality, integrity and availability.

---

<sup>4</sup>Indeed, it has been estimated that the building costs of a house sum up to about 20% of the total of all costs of living in the house[Kol04]. Adding 10% or even 20% to the building costs due to automation will be an investment that will pay itself back fairly quickly, as well as add to the resale value of the house.

**The third and final interview** touched a whole slew of security topics, questions as well as answers, risks as well as chances. “Computer security is like a sewer – nobody notices it until it goes bad” started the interview session. People see information security as an image, a brand – something to trust in – and for the consumer, trust is the most important value to uphold. It is even more important than security itself. It is less a problem for the user if she is informed that there has been an incident, which systems are affected, and that normality will be restored by a clearly given time – and that the service providers stick to that time – than that something somewhere has gone terribly wrong and the service provider isn’t informing her. People must feel that the security management process is predictable. Things are bad now but they will be taken care of by someone (or something) competent. The author thus realized that this is in fact very much like management of large crowds in a crisis: keep the crowd informed even of bad news, tell them when there is going to be change, and keep your promises [UNO03]. If crisis/security management has a good image, people will not be too concerned about the crisis/security, they consider that the crisis/security issue is being taken care of. This is what happens with security software such as anti-virus products and firewalls when they work properly. For the society at large, the only thing that matters is long-term stability, which implies that the user is on top of what is happening, and that she feels there is somebody who takes care of issues, in this case security issues, for her.

One important issue with the future home network is how do we establish a securely interoperating environment (this implies that the huge task of creating just an *interoperating* network is already handled). Kari suggests a trust model and process, reminiscent of the Resurrecting duckling security model [SA99] (see § 5.4), and a very simple user interface where even the most mundane appliance such as the proverbial Internet Toaster would have a Reset button and three LEDs: a green to indicate that everything’s okay, a yellow one to indicate that there is an issue that the user should investigate and a red one to mark that the device is off the net. When the new appliance is brought into the home network, it will holler for a parent device and imprint itself on that device. In contrast with Stajano’s Resurrecting duckling model, Kari suggests that often this process can be made completely without user interaction – if the device happens to imprint itself on the neighbour’s network, the user would just press the reset button and the imprinting process would restart. The same reset button would be used to dissociate the device from the home network if, for example, the user wanted to sell or give the device away. Of course, not all devices should be easily dissociateable; a cell phone, laptop computer, digital camera or portable music player should firmly hold on to their imprinting to be useless to thieves.

A very practical way of authenticating oneself to a device or imprinting

it, as suggested by Kari, is by touch. If the user wears an authentication token, say a wrist watch or some kind of jewelry with galvanic contact to the user, the user's skin could act as transmission medium between the authentication token and the device. After an initial 'handshake', the rest of the authentication process could happen over the air. The authenticated session ends automatically when the fob exits the radio proximity.

Such a fob infrastructure of course implies that the user carries his or her authentication token at all times, or, that an authentication token can be made that the user wants to carry at all times. Authentication to the token itself could happen with a fingerprint, since the user's fingers will be employed to put on the token. If the token is taken off, or drops off the user, the token would go into a state where it doesn't act as an authentication agent for the user, i.e. go into standby mode. A model like this would afford a very unobtrusive way of authenticating to any device, like a computer terminal or a phone: just sit down by any computer or pick up any phone, and it's "yours". If the authentication token is compromised, a requirement is that the user's chain of trust can be rebuilt even on top of the compromised chain. The previous chain needs to be revoked and a new one restored on the top of it.

Different kinds of devices benefit from different methods of imprinting. A personal item such as a new computer demands a higher degree of security and certainty in the imprinting process and would be imprinted by touch. On the other hand, with the newly bought Internet Toaster, it is enough just to plug it in and press the reset button. If it shows up on a monitoring station, the appliance has successfully found a 'mother duck', if not, just press the reset button anew.

Guests of the home would have similar authentication fobs. Thus, when a guest comes around, we first associate (or imprint) the guest, through her authentication token, with our house, by which, through a chain of trust, our guest's devices are also considered 'friends' of the house and allowed certain access to the services the house offers. These services offered would of course be contextually dependent on for example whether an inhabitant is around; if the house is 'alone', it would enter 'a state of higher paranoia'. To maintain security, no identity fobs would answer a beacon that it doesn't have a trust relationship with.

Technically, all this is almost trivial. In practise it is complicated, but more of logistic and bureaucratic reasons, and most complicated is the management of rights. As mentioned in the interview with Risto Linturi, different users of the home have different roles, but also the same users of the home have different roles, and it is not always evident which role a user appropriates for a given task. Role management is a part of the rights management process, and it is a complex one.

Privacy is contextually dependent, and privacy measures need to include a controlled override mechanism. People tend to have two “levels” of privacy need, normality and emergency. For example, in normality, a user would definitely want to keep her medical information private, but in a case of medical emergency, it may be critical to disclose this information to a trustworthy party, e.g. the medical personnel. This transaction also needs to be logged for traceability.

There is a future business in security. A consumer could buy himself a certain guaranteed level of security (as we have stated before there is no absolute security; see § 4.3). In this way, the user could be allowed to bring new and exotic devices such as the Internet Toaster to his home and the security service will check that the appliance is compatible with the rest of the user’s home environment and monitor that it behaves. Similarly, trust is a future business area. Trust can be brokered, so that a consumer does not have to trust a merchant but can pay a service provider to issue and broker the trust, like an issuer of a validation stamp or a seal of approval. In a way, the trust broker then is in the insurance business; if the merchant is not trustworthy, the trust broker takes the hit. In another sense, this is an insurance against stupidity, but the question is, do we really want a guard against stupidity? People must understand not to trust blindly in security mechanisms, for example if one is given a 100 euro bill that is pink, one shouldn’t go and accept it as valid currency even if the bill verification system accepted it.

The security service could certainly also act proactively; if there are atypical patterns in the home network traffic, for instance, the monitoring service of the security provider could inform the user that something might be unwell.



# Chapter 7

## Analysis

In this chapter we analyze how the users employed access control in their homes. We note that in substantial measures, neither conflicts nor overlaps between the security views of home users and security experts were found. We present some ideas on how to apply the knowledge gained in this thesis on home network access control and round off the chapter with some topics for further research.

### 7.1 Family interview findings

Manifestations of many access control methods were found during the family interviews. The most important one was that access control is governed by social conventions or ‘the social barrier’. Some things are considered private or off-limits and the respect of this privacy extended seamlessly from the physical world to the digital domain. The ‘Big Stick principle’ [Sta02b] was also frequently manifested with portable digital devices, both devices seen as personal (music players) or common/shared (the children’s digital camera).

The most relevant finding was that the people worked in a clearly GRBAC-like mindset. Most notably, the parents had different roles for different tasks. At one time, the parent could be doing her work from home, at another time she was organizing her children’s hobby activities, or her own freetime engagements. Sometimes the parent took the role of the systems administrator. To a certain degree, there were computers that also would serve as work and study tools as well as entertainment machines, but usually computers meant for work were used strictly for work. The children would of course have different roles to their parents, but they themselves didn’t show as strong a distinction between their ‘student’ and ‘leisure time’ roles, in part because their student and leisure time activities largely involved the same friends. In a similar fashion, the teenagers often could use instant messaging

in parallel with studying activities, regardless if the communication was part of the studies or not.

People with security knowledge usually object against employing ‘security by obscurity’. This method was successfully (albeit unwittingly) employed by a family sharing their photographs with their friends and extended family using identifiers (URLs) and metadata (descriptions such as names) that could not be easily linked to the family. As a test, we tried to find the family’s photos using a few Internet search engines (Google, Yahoo!, MSN<sup>1</sup>), but did not manage to come up with anything. Contrawise, the author’s photos were easily found from a photo sharing service using those same search engines. It can be concluded that with current search technology (and our search methodology), security by obscurity works. Nevertheless, with all our security know-how, we still advise against the security by obscurity approach.

Users wanted technology to be simple<sup>2</sup>, be invisible, and work. One parent told us that she got very stressed about even the thought of technology not working, and doing so in indecipherable ways, and if there was something she did not need right now it was more stress. Another manifestation that simplicity is a desired property came from the Mac using family who were happy to use Macs because “they just work”.

A recurring request, even a requirement, from the users is that the home network must be easy to use. It must be unobtrusive and not get in the way. This is clearly in line with both usability literature in general and HCIsec guidelines. People are not very interested in maintaining their home network’s security. What we didn’t see were the demands that the system should be transparent, i.e. keep the user informed on its state and what it is doing. This is probably due to the focus of the user interviews – we only wanted to see what the users might do with home networking technology, not test some technology we had built.

To respond to these findings, we can conclude that access control needs be easy to maintain and it should be flexible; it’s hard to predict who is going to need what rights. When a guest arrives, he must easily be allowed proper access. In short terms, network security should be built to be a convenience factor as well as a security factor. After all, the home network should ultimately make the inhabitants’ life easier. Once properly deployed, things on the home network will seamlessly do what the user intends. Convenient security is after all one key goal for HCIsec. Role-based access control, and in particular, Generalized role-based access control seems to model the dynamics of home rather well and the author believes it would be beneficial to

---

<sup>1</sup>MSN search is now branded Live search

<sup>2</sup>In this context, we use the word *simple* in its most flattering, and designer-demanding way; not as feature-restrictive but as un-complex, yet empowering.

base the home network's infrastructure on this model.

## 7.2 Comparing family and expert opinions

Valuable input on security issues were collected from the expert interviews even if it all could not be applied to the focus of this thesis. The interview with Ben Lavender circled around the distribution of media to home users, the Risto Linturi interview about the importance of the infrastructure. The Hannu H. Kari interview provided food for thought to the home network ecosystem, as well as some practical ideas on what to put in it. The family interviews on the other hand gave an insight in how actual people see their environments and how their interactions with it functioned.

There were no immediate conflicts between the views of the home users and those of the security experts. Rather, they were complementing each other and apart from the fact that all said the security and the home network must be easy to use, there was little overlap.

## 7.3 A sketch for a usable access control ecosystem

A home network security design and control system should *guide* the user into making sound, secure and informed choices and allow the user enough freedom to make the environment their own. But how much control should the user have? A common conception, especially among programmers of a hacker mindset, is that more user control leads to a better use experience and hence, is always better. An 'intelligent' product that makes all the choices for the user without giving any real options or telling her why, will frustrate the user and make her trust the product less (even if the decision were the same that the user would have picked, given the option of an informed choice) [Nor07]. But while user centred design puts the user in control of everything, this does not mean that she should necessarily *have* to control the specifics of the user experience. In fact, when graphed, a user satisfaction to level of control would look bell-shaped (Figure 7.1 by Kathy Sierra, used with permission); giving a user too little freedom and control can be just as frustrating as allowing the user to tinker with each available detail [Sch04b]<sup>3</sup>. The challenge is to find how to present the user with just the right amount of user control.

The geometry of this supposed bell shaped curve is of course different for each user. The more experienced the user is, the more she will enjoy hav-

---

<sup>3</sup>The discussion herein really must be credited Kathy Sierra, who in a wholly unscientific article described this thought, along with the vaguely bell-shaped curve in [http://headrush.typepad.com/creating\\_passionate\\_users/2007/02/how\\_much\\_control.html](http://headrush.typepad.com/creating_passionate_users/2007/02/how_much_control.html)




Figure 7.1: A hypothesis: User control versus user happiness

ing the possibility to deal with each detail in the product. As an extreme example, a professional driver would sometimes switch off the automatics of their car since the automation will not allow them to deploy their skills. Still, based on both the literature and interviewes presented in this thesis, the typical user of the home network should foremost just describe *what* the thing should do, and unless told otherwise, leave the “how”-part to the thing itself.

We can also assume that the security system can be built with a greater knowlege in security requirements than the average user possesses. Thus, a requirement for an access control system on a home network is that the user describes the intent of the security environment, and the environment, in a dialogue with the user, builds up the security policy for itself. A good starting point would be some sensible default scenarios. An important aspect is that the system should be able to explain why it suggests such and such a solution, so that the user can make an informed decision whether this one is good for their environment. The system would continuously test itself for security “leaks” and if need arises, the system would strongly recommend against solutions that are insecure in the policy.

From the family interviews we could see clear indications of a Role-based access control (RBAC) thinking in place. While the families did not talk in GRBAC terms, we can see that they would implicitly work with GRBAC concepts. The user roles naturally mapped to the family hierarchy; the parents were the Deciders, with the other parent also taking the Expert (or system administrator) role. The children had less privileges on the network than their parents, and the privilege varied with the technical expertise and age of the children. The same was true for devices; depending on who was using it (or, on another axis, in what context), a computer could either be one for entertainment or work, and have different capabilities. For example, in one family, the childrens’ computers Internet access was regulated by a time-based firewall rule which in effect stated “no Internet after bedtime”.

### 7.3. A SKETCH FOR A USABLE ACCESS CONTROL ECOSYSTEM<sup>73</sup>

A security interface to assist the family in creating an access control policy – and more broadly, a security policy – could be based on Generalized role-based access control (GRBAC) concepts. Such a security interface would need sensible suggestions for different household configurations, listing inhabitants, peer groups and so on as default (and configurable) user ‘roles’. The user roles would be hierarchical, eg. so that the parent role would be part of the inhabitant role. The creation and maintenance of the security policy need in no way be a solitary activity of the Home network systems administrator. To promote transparency of the policy, its creation can very well be a family activity. As a heritage from Discretionary access control (DAC), inhabitants would be allowed to classify resources they own, i.e. resources in their subtrees. Ownership of such physical resources may be signified using an application of The Resurrecting Duckling model [SA99].

While the user would, assisted by the software, need to think of all people of the home network and their roles, the home network infrastructure in a fully networked home should be able to discover all devices and services within. This is one of the basic functionalities of UPnP and other service discovery protocols discussed in § 3.6. The security interface would help the administrative user in categorizing the devices, their use and their audience. A security interface would then assist the user in describing different situations, the family is home and away, it is daytime (in subjective or chronometric terms), night time, party time, barbecue time on the backyard and so on. The situations can then be used to describe how the home network should behave in those different states.

We envision that the full power of a GRBAC architecture need not be unleashed for it to be useful home use. In a full GRBAC architecture, a user could have different, even conflicting roles at different times, but the most likely change of role a user would need is that between a parent and an administrator. The administrator role should be reserved for administrative jobs, and discourage the user performing their ordinary activities as administrators. A similar approach is taken in Windows Vista and many modern Linux systems: if the user needs elevated privileges for an administrative task, the user would authenticate herself as a member of the administrators group and do the administrative work “on top” of their ordinary session, not start a new login session as administrator, where the user may get too comfortable for system security.

What identification and authentication system would be needed on a home network system depends on what is to be authenticated and in which context. There are three types of devices: the devices that the users own, which are tied to the home network (authenticated), devices of known users (authenticated) and unknown devices (unauthenticated). Using provisions from or similar to UPnP, device authorization of authenticated devices can

be largely automated, as can discovery of unauthenticated devices. In a similar fashion, a user can be in his home, in another home which is interconnected to the user's home, or at a public site. All three scenarios give different demands to security [SV04]. Mobile devices would need to make an assumption, or consult the user, on their context, to assume a suitable "level of paranoia". What type of user identification and authentication means is to be used also depends on context. Given the technology, users themselves would most unobtrusively be identified using biometric methods. The identification and authentication backline would be augmented by state machines, Markov chains and other means of learning the habits of its inhabitants. Where the level of protection (and confidence) allows, the ever-present social barrier (and an audit trail) would mean that a "good guess" from the system is enough. Such fuzzy logic could support the scenario that anyone may open the medicine (or bar) cabinet as long as at least one adult is home. We now present a few ideas that may be applicable for increasing security on the home network, in a user-centred manner.

### 7.3.1 The Door: responsibility driven access control

The home network environment is a complicated one with many needs for access control and great need for flexibility. Still, the most commonly exhibited form of access control was that of self-imposed ethical and moral standards: stuff that was somebody else's private property was left alone. While ethics may be a rather complicated thing to describe to a computer, we can get by all the complexity by simply ignoring it, or rather, delegating the understanding of what's right and wrong to the users, and leave the computationally easy parts like logging to the machinery.

To reflect this reality, a 'Door' metaphor is proposed as a means for access control between humans and the resources they are trying to access. A user could choose whether to breach the requested level of privacy, so the system really builds on trust between the users. Essentially, the Door is like Stajano's Big Stick, except that its user is expected to decide whether they are going to use it for other reasons than just the fact that they are able to.

A door would need to identify and audit, but not to authenticate. In fact, a naïvely trusting door would not even do identification or the auditing. It would just be like an ordinary door with a "Do not enter" sign on it. Like physical doors, the metaphorical ones can be closed or left open, and they can be locked. Like some doors, they have a label with their owner's name on it, an information that is conveyed to the human who is considering walking through the door. Like any unlocked, unguarded door, this one too will let anyone through that opens it. And unlike most doors, the ever watchful digital one logs all accesses through it. A digital door like the

### 7.3. A SKETCH FOR A USABLE ACCESS CONTROL ECOSYSTEM<sup>75</sup>

one proposed would build on the built-in ability of most people to choose between right and wrong. The access control would be self imposed and driven by responsibility.

A Door would be appropriate in a multi-level security setting, like in a home environment. As an example, an inhabitant would need to be regarded as an adult to access the medicine- or bar cabinet, and they would need to be the owner of the single malt whisky bottles within to access those. Accesses will be logged by camera, and it is up to the application, or bar/medicine cabinet owner to look at the logs. Such a system would promote responsible behaviour in the home network environment, as well as greatly simplify the creation, maintenance and, not the least, enforcing the access control on the home network.

Doors could only be used where the people involved in the access transaction can be trusted, or when the result of a breach either isn't catastrophic or is easily reversible. Undesired external principals need still to be firewalled out. The latter situation, Doors open for everybody, can be seen in the philosophy of Wikipedia<sup>4</sup>, "The encyclopedia anyone can edit." The openness is a centerpiece of the philosophy of the Wikipedia. As anyone is welcome to contribute, many will, and the works of vandals can easily be reversed by a rollback. While it was assumed above that ethics aren't easily taught to a machine, Doors might be beneficial also for machine-to-machine communication; a service might answer back that it won't have a problem with accessing a request larger or more demanding than a certain number and it is up to the requesting machine to decide on how and whether to offload the request.

A Door would both do away with the authorization part of the access control, and build upon the social barrier and respect for privacy within the home. Of course, if we wanted good auditing built in, the Door would still need to be able to perform identification of the user. A hybrid digital-non-digital Door would be a physical door equipped with a video camera; when the user opens the door, he will get his photograph taken, logged and timestamped, and the owner of the Door would be able to see the log at a later state.

The Door could also be used in emergency context, if a person would be in dire need of emergency medical attention, a helper could open the door as a 'manual override' to private medical information. This would be modern-world equivalent of the 'SOS-passi' bracelets and pendants that were fashionable in Finland in the 1980s and early 1990s, and contained a folded strip of paper with 'emergency medical information'. If the individual was in a conscious state, the information would easily be kept private. Indeed an implementation of this already exists: many people have marked the person to be contacted In Case of Emergency with the abbreviation 'ICE' on their

---

<sup>4</sup><http://www.wikipedia.org>

cell phones' directories (phone books). The cell phone is a personal artifact, and the phone directory even more so, but in an extraordinary situation the privacy is meant to be breached.

### 7.3.2 Other security ideas

Many network technologies work so that they listen to all incoming data and discard the traffic that was not intended for them. In a wireless ad-hoc environment, traffic is often encrypted for privacy, which can be hard on 'peanut devices' with low power and computing resources. For peanut devices, public key cryptography, being an approximately three orders of magnitude more computationally intensive job than a symmetric cryptography approach of comparable strength, is generally out of the question [Sta02a] (and cryptography is generally something that cannot be delegated). Packet level authentication (PLA) has been suggested as a solution for many security problems in military grade network environments [CLK05], but a scaled-down solution of PLA employing a lightweight encryption scheme could be applied for home networks. Thus, the encryption of traffic could serve as a form of addressing or packet selection at the receiving side, in addition to providing access control and resiliency against eavesdropping and denial of service attacks. If the node can open the packet, i.e. the contents of the decrypted data is well-formed/readable, it is of relevance and the payload will be inspected; otherwise it is discarded. This method also borrows from two-way authentication and "identify friend or foe" applications.<sup>5</sup>

In a wireless distributed network, it is of course not possible to deny a device physical access to the network itself. However, it is quite possible to do the same thing in a distributed fashion; each device can make an individual decision not to communicate with a device that does not fulfill certain security criteria.

In an adaption of the usability guideline that the correct usage of a thing should be the most usable one, also the most secure actions on a thing should be the easiest and most natural one to take. To extrapolate on the thought a bit further, it should be hard to make something insecure. It is debateable though if a user should be denied to completely destroy his own security (as long as it has no implications on his surroundings). The author opines that a user should be allowed "to shoot himself in the foot", but the system must make it absolutely clear to the user what the consequences are that this is not a good idea. A home network security system or component should come with sensible defaults, a good set of security templates or patterns, and be able to configure itself by ask only a few relevant questions, like "what is

---

<sup>5</sup>It is possible that this proposal re-invents an already invented wheel, but the author was not able to identify previous art in this context.

your level of paranoia?”

## 7.4 The applicability of Contextual design

In the research phase of this work, we used parts of Contextual design, namely the Inquiry, Interpretation and Consolidation steps. As our research material was relatively small, and the same core team was present at all the interviews, there was little need for the modelling phase of Contextual design. Since the interviews were done to create an understanding of our users, not a product for them, the Work redesign, User environment design, Prototyping and Implementation phases were left out of the exercise. Given the interviews, materials and especially the excellently co-operative families, these steps could be incorporated in future research if so needed.

So why was Contextual design employed? Indeed, why use Contextual design at all, anywhere? Because the only way to figure out what users need is to ask them, and a well-documented method of doing so is embodied in the Contextual design process. While all of Contextual design was not applicable to this work, we used the parts that were relevant.

Even though we only used a ‘light-weight’ application of Contextual design, the exercise as immensely useful in getting real data “from the field”. For our purposes, the unadulterated Contextual design process would have been overkill both for the research but not the least for our users, who we did not want to bother overly with our presence.

We discussed using additional tools to gather more information from our user. The tools considered were logbooks or diaries, questionnaires, cultural probes<sup>6</sup> and asking the teens to take pictures with their cell phones and sharing them with us (which in itself is a light-weight diary or cultural probe). While all these means would have brought more input to the work and the project, and would have been interesting to use, we opted to do without them for two simple reasons. We did not have the resources to execute the work needed, and we did not consider that gathering this information would have yielded enough a benefit to justify the work, within the scope of this thesis and where we were at the time with the InHoNets research work.

## 7.5 For future research

One issue, hinted at in chapter 3, was that the same people are different kinds of users when in different environments. We know from before that on a home network, users are even more goal oriented than while on an

---

<sup>6</sup><http://www.infodesign.com.au/ftp/CulturalProbes.pdf>

office network. It could be interesting to research how the expectations and attitudes towards technology really differ, and why, between these environments. Will a skilled IT professional turn into an easily frustrated neophyte at home?

A burning question that taunted the author during all this work was “what’s wrong with the users” when they do not care about security. In reality of course nothing is wrong with the users and it is just a question of what is important to the user and what is understandable. Since normal people do have a sense of what real-life security is of importance and what should be done to guard oneself against the hostilities of everyday life, it should not be impossible to find a way to communicate home network security to a user of the technology in a compelling way.

**Advanced GRBAC on the home network** As mentioned in § 4.4, the Role-based access control (RBAC) model exists as basic RBAC, with the extensions  $RBAC_1$ ,  $RBAC_2$  and  $RBAC_3$ .  $RBAC_1$  allows inheritance of roles,  $RBAC_2$  adds constraints to roles (i.e. some combination of roles are unacceptable, c.f. the Chinese Wall security policy [BN89]), and  $RBAC_3$  is the union of RBAC levels 1 and 2. While inheritance is a part of Generalized role-based access control (i.e. GRBAC is really an extension of  $RBAC_1$ , not “ $RBAC_0$ ”) the constrained GRBAC has not yet been proposed.

While it would be trivial to propose  $GRBAC_2$  as a Constrained generalized role-based access model per the descriptions above, this would in the context of this thesis be an exercise in worthless trivia, unless a compelling reason to apply constraints in a home network environment can be identified. This exercise will be left to future research.

## Chapter 8

# Conclusion

The home network as seen in this thesis is an early form of the ubiquitous computing network, an environment with a large amount of small specialized computing devices working together with users who do not want to perceive the environment as a computing environment. This kind of a home network does not yet exist, and neither are the usability and security aspects of it clearly outlined. Therefore, we have an excellent opportunity to consider the human-computer interfacing (HCI) and security of the home network together as HCIsec, and be prepared in time before such networks become commonplace. Indeed, security and a good user experience are essential factors for making home networks come true.

In this thesis, aspects of the home network have been presented, as well as security solutions that when adapted to the home network environment should be usable there. Both existing literature, security professionals and real families have been heard. In particular, a security policy builder based on Generalized role-based access control models has been suggested, due to the fact that the families displayed GRBAC-like thinking in their views on access control in their homes. The policy builder would guide the user into building a secure environment, based on familiar concepts and sensible defaults.

Another pattern that seemed to govern the internal access control behaviour was the ‘social barrier’ or simply social (and territorial) conventions. Due to this, it would in cases where a higher level of security isn’t needed be appropriate just to mark a resource as private with a ‘door’ metaphor. We often saw inhabitants of the families keeping their personal files in folders, readable and accessible to anyone, only “protected” with their name on it; the social convention dictates that you do not look inside another’s folder without permission. This would make access control for users within the family a whole lot simpler. Finally, the inhabitants want to spend as little effort on maintaining the home network and find it as easy as possible to

enjoy the possibilities the network provides. The Door access control should only be used when the ‘protected’ asset can handle a “Door breach”

Current realizations of domestic technology hardly exhibit features desirable for a seamless home network environment. They do not interoperate well, they are not designed with multiple users in mind – each with different aptitude, goals and even roles. Products are hard to both install, extend and operate, even for the seasoned technology enthusiast. Even a simple thing like watching digital photos on a television can be quite a daunting task.

On a home network, devices will need to virtually install themselves and operate without a central security server. There are solutions on different levels to this problem. First, the Resurrecting Duckling model, especially in its extended, ad-hoc installation, would be very useful for approving devices to the home network. While there may not be any realizations of the Duckling model yet, UPnP is a worthy family of protocols to organize the backbone of the home network. Implementations already exist, though UPnP security realizations do not seem to have escaped the lab environments just yet.

A home network should be aware of its users’ whereabouts, desires and intents and be able to respond in accordance with the principle of least surprise. For this, the home network needs to be able to sense who is doing what and where, and put this information into context. It would also need to convey its intents and status information in an unobtrusive manner to the inhabitants, in order for the inhabitants to feel (as well as be) they are in control of the home and not the other way around. Sensors for at least rudimentary guesses for most of these questions already exist in the form of ‘smart floors’, ultrasound sensors, motion detectors, vision-based sensors, as well as calendars, clocks and mathematical algorithms spanning Markov chains, fuzzy logic and artificial intelligence. But they do not yet exist on the mass market in the size, price and capability factors envisioned in this text. While the consumer can not yet go to the consumer electronics store and buy a bag of peanut devices, a rough sketch of the functionality of this future ubiquitous home network can already be built. And indeed, is already being built by early adopters, researchers and enthusiasts. The future is already here. It’s just not evenly distributed yet.

# Bibliography

- [Ada84] Douglas Noël Adams. *So Long, and Thanks For All the Fish*. PAN books, 1984. ISBN 0345391837. 30
- [All04] Digital Living Network Alliance. Overview and white paper. [http://www.dlna.org/about/DLNA\\_Overview.pdf](http://www.dlna.org/about/DLNA_Overview.pdf), June 2004. 9, 18
- [And01] Ross Anderson. *Security Engineering*. Wiley, 2001. ISBN 0-471-38922-6. 36
- [Bah02] Anna Bahney. High tech, under the skin. New York Times, 2. Feb 2002. [http://amal.net/blog/links/2006-02-02\\_-\\_High\\_Tech\\_Under\\_the\\_Skin\\_-\\_New\\_York\\_Times.pdf](http://amal.net/blog/links/2006-02-02_-_High_Tech_Under_the_Skin_-_New_York_Times.pdf). 41, 51
- [Bar05] Jakob E. Bardram. The trouble with login: on usability and computer security in ubiquitous computing. *Personal Ubiquitous Computing*, 9(6):357–367, 2005. <http://www.springerlink.com/content/h1744q1332032788/fulltext.pdf>. 50
- [BG02] Gordon Bell and Jim Gemmell. A call for the home media network. *Communications of the ACM*, 45(7):71–75, 2002. <http://doi.acm.org/10.1145/514236.514237>. 9
- [BH97] Hugh Beyer and Karen Holtzblatt. *Contextual Design - Defining Customer-Centered Systems*. Number ISBN 1558604111. Morgan Kaufmann, 1997. 30, 58
- [BH99] Hugh Beyer and Karen Holtzblatt. Contextual design (article). *Interactions*, 6(1):32–42, January-February 1999. <http://delivery.acm.org/10.1145/300000/291229/p32-beyer.pdf?key1=291229&key2=2942085511&coll=GUIDE&dl=GUIDE&CFID=1545619&CFTOKEN=77022284>. 57
- [BN89] Brewer and Nash. The chinese wall security policy. In *Symposium on Security and Privacy*, pages 206–214, Oakland, CA, USA, May 1989. IEEE. <http://www.cs.nmt.edu/~doshin/t/s06/cs589/pub/4.BN-ChWall.pdf>. 37, 78

- [CAJ03] Lynne Coventry, Antonella De Angeli, and Graham Johnson. Honest it's me! self service verification. In *Workshop on Human-Computer Interaction and Security Systems*. ACM, April 2003. 41, 42
- [CAMN<sup>+</sup>02] Roy H. Campbell, Jalal Al-Muhtadi, Prasad Naldurg, Geetanjali Sampemane, and M. Dennis Mickunas. Towards security and privacy for pervasive computing. In Mitsuhiro Okada, Benjamin C. Pierce, Andre Scedrov, Hideyuki Tokuda, and Akinori Yonezawa, editors, *ISSS*, volume 2609 of *Lecture Notes in Computer Science*, pages 1–15. Springer, 2002. <http://srg.cs.uiuc.edu/gaia/papers/towards-percomp-security.pdf>. 18, 43, 49
- [CER01] CERT. Home network security. [http://www.cert.org/tech\\_tips/home\\_networks.html](http://www.cert.org/tech_tips/home_networks.html), 2001. 33, 47
- [CGR04] Domenico Cotroneo, Almerindo Graziano, and Stefano Russo. Security requirements in service oriented architectures for ubiquitous computing. In *MPAC '04: Proceedings of the 2nd workshop on Middleware for pervasive and ad-hoc computing*, pages 172–177, New York, NY, USA, 2004. ACM Press. 25, 53
- [Cla94] Roger Clarke. Human identification in information systems: Management challenges and public policy issues. *Information Technology & People*, 7(4):6–37, December 1994. <http://www.anu.edu.au/people/Roger.Clarke/DV/HumanID>. 42
- [CLK05] Catharina Candolin, Janne Lundberg, and Hannu H Kari. Packet level authenticatoin in military networks. [http://www.tcs.hut.fi/~hkh/cv/pubs/AIWITSC2005\\_Kari\\_et\\_al.pdf](http://www.tcs.hut.fi/~hkh/cv/pubs/AIWITSC2005_Kari_et_al.pdf), 2005. 76
- [CMA00] Michael J. Covington, Matthew J. Moyer, and Mustaque Ahamad. Generalized role-based access control for securing future applications, November 2000. <http://citeseer.ist.psu.edu/493143.html>. 55
- [CN02] Mark D. Corner and Brian D. Noble. Zero-interaction authentication. In *MobiCom '02: Proceedings of the 8th annual international conference on Mobile computing and networking*, pages 1–11, New York, NY, USA, 2002. ACM Press. 51
- [Den76] Dorothy E. Denning. A lattice model of secure information flow. *Commun. ACM*, 19(5):236–243, 1976. 15
- [DGdlFJ04] Paul Dourish, Rebecca E. Grinter, Jessica Delgado de la Flor, and Melissa Joseph. Security in the wild: user strategies for

- managing security as an everyday, practical problem. *Personal and Ubiquitous Computing*, 8(6):391–401, 2004. 44
- [Dha00] Rachna Dhamija. Hash visualization in user authentication. In *Short Paper Proceedings of the Conference on Human Factors in Computing Systems*. CHI2000, April 2000. 50
- [DP00] Rachna Dhamija and Adrian Perrig. Déjà vu: A user study using images for authentication. In *Proceedings of the 9th USENIX Security Symposium*. USENIX, August 2000. <http://people.deas.harvard.edu/~rachna/papers/usenix.pdf>. 50
- [EKO<sup>+</sup>99] Irfan A. Essa, Cory D. Kidd, Robert J. Orr, Gregory D. Abowd, Christopher G. Atkeson, Blair MacIntyre, Elizabeth Mynatt, Thad E. Starner, and Wendy Newstetter. The aware home: A living laboratory for ubiquitous computing research. In *Proceedings of the Second International Workshop on Cooperative Buildings - CoBuild'99*. Georgia Institute of Technology, October 1999. [http://www.awarehome.gatech.edu/publications/cobuild99\\_final.PDF](http://www.awarehome.gatech.edu/publications/cobuild99_final.PDF). 17
- [Ell02] Carl Ellison. Home network security. *Intel Technology Journal*, 06(04), November 2002. [ftp://download.intel.com/technology/itj/2002/volume06issue04/art04\\_security/vol6iss4\\_art04.pdf](ftp://download.intel.com/technology/itj/2002/volume06issue04/art04_security/vol6iss4_art04.pdf). 1, 6, 19, 23, 38, 48, 57, 58
- [Ell03] Carl Ellison. *UPnP<sup>TM</sup> Security Ceremonies design document*. UPnP Forum, version 1.0 edition, 3 October 2003. [http://www.upnp.org/download/standardizeddcp/UPnPSecurityCeremonies\\_1\\_0secure.pdf](http://www.upnp.org/download/standardizeddcp/UPnPSecurityCeremonies_1_0secure.pdf). 21
- [FK92] D. Ferraiolo and R. Kuhn. Role-based access controls. In *15th NIST-NCSC National Computer Security Conference*, pages 554–563, 1992. 37, 54
- [FSH03] Ivan Flechais, M. Angela Sasse, and Stephen M. V. Hailes. Bringing security home: a process for developing secure and usable systems. In *NSPW '03: Proceedings of the 2003 workshop on New security paradigms*, pages 49–57, New York, NY, USA, 2003. ACM Press. 35, 45
- [GEND05] Rebecca E. Grinter, W Keith Edwards, Mark W Newman, and Nicholas Ducheneaut. The work to make a home network work. In H. Gellersen et. al., editor, *Proceedings of the Ninth European Conference on Computer-Supported Cooperative Work*, pages 469–488, Paris, France, September 2005. EC-

- SCW. <http://www2.parc.com/csl/members/nicolas/documents/ECSCW05.pdf>. 1, 6, 9
- [Gib93] Willam Gibson. Interview on “fresh air”. NPR National Public Radio, 1993. <http://www.npr.org/templates/run downs/rundown.php?prgId=13&prgDate=31-Aug-1993>. 9
- [GL85] John D. Gould and Clayton Lewis. Designing for usability: key principles and what designers think. *Commun. ACM*, 28(3):300–311, 1985. <http://www.research.ibm.com/comp sci/spotlight/hci/p300-gould.pdf>. 30
- [Gre06] Adam Greenfield. *Everyware: The Dawning Age of Ubiquitous Computing*. Peachpit Press, 1st edition edition, 10 March 2006. ISBN: 0321384016. 16
- [Gro04] Elisheva F. Gross. Adolescent internet use: What we expect, what teens report. *Applied Developmental Psychology*, (25):633–649, 2004. <http://www.center-school.org/pko/documents/AdolescentInternetusepdf.pdf>. 60
- [Gru92] J. Grudin. Utility and usability: Research issues and development contexts. *Interacting with Computers*, 4(2):209–217, August 1992. I haven’t read this yet, but Nielsen refers to it. 28
- [Hol99] Ursula Holmström. User-centered design of security software. In *Human factors in Telecommunications*, Copenhagen, Denmark, 4.–5. May 1999. <http://www.tml.hut.fi/Research/TeSSA/Papers/Holmstrom/hft99.pdf>. 45
- [HSU04] Dieter Hutter, Werner Stephan, and Markus Ullmann. *Security in Pervasive Computing: First International Conference, Boppard, Germany, March 12-14, 2003. Revised Papers*, volume 2802 / 2004, chapter Security and Privacy in Pervasive Computing State of the Art and Future Directions, pages 285 – 289. Springer Berlin / Heidelberg, 2004. ISSN: 0302-9743, ISBN: 3-540-20887-9, DOI: 10.1007/b95124, <http://www.springerlink.com/content/x15cff7t25wgv6cn/fulltext.pdf>. 49
- [HT04] David M. Hilbert and Jonathan Trevor. Personalizing shared ubiquitous devices. *interactions*, 11(3):34–43, 2004. 40
- [HWW05] Karen Holtzblatt, Jessamyn Burns Wendell, and Shelley Wood. *Rapid Contextual Design, A How-To Guide to Key Techniques for User-Centered Design*. Number ISBN 0-12-354051-8. Elsevier/Morgan Kaufmann, 2005. 32

- [inh06] Interconnected broadband home networks research project. <http://www.tml.hut.fi/Research/inhonets>, 2005–2006. 2, 8, 19, 20, 57
- [ISO98] ISO (International Standards Organization). *Guidance on Usability*, 1998. ISO 9241-11. 28
- [IT03] ITU-T. Generation and registration of universally unique identifiers (uuids) and their use as asn.1 object identifier components. Recommendation X.667, ITU-T, 13 September 2003. <http://www.itu.int/ITU-T/studygroups/com17/oid/X.667-E.pdf>. x
- [JŘ03] Václav (Vashek) Matyáš Jr. and Zdeněk Říha. Toward reliable user authentication through biometrics. *IEEE Security & Privacy*, 1540-7993/03:45–49, 2003. 43
- [JSL<sup>+</sup>04] R Jimeno, Z Salvador, A Lafuente, M Larrea, and A Uribarren. An architecture for the personalized control of domotic resources. Poster article at EUSAI 2004, November 2004. Eindhoven, the Netherlands. 40
- [Kar06] Hannu H. Kari. Personal interview, 14 August 2006. 18, 51, 53
- [KLKY02] Dong-Sung Kim, Jae-Min Lee, Wook Hyun Kwon, and In Kwan Yuh. Design and implementation of home network systems using upnp middleware for networked appliances. *IEEE Transactions on Consumer Electronics*, 48(4):963–972, Nov 2002. ISSN: 0098-3063, Digital Object Identifier: 10.1109/TCE.2003.1196427. 9, 24
- [Kol04] Tero Kollanus. Smartpirtti, 2004. <http://www.nettitalo.com/smartpirtti/Smartpirtti-inssityo.pdf>. 64
- [LMS05] P. Leach, M. Mealling, and R. Saltz. A universally unique identifier (uuid) urn namespace. Technical Report RFC4122, Network Working Group, July 2005. <http://tools.ietf.org/html/rfc4122>. x
- [MA01] Matthew J. Moyer and Mustaque Ahamad. Generalized role-based access control. *icdcs*, 00:0391, 2001. 54
- [Mas43] Abraham Maslow. A theory of human motivation. *Psychological Review*, 50:370–396, 1943. 17
- [May06] R. Mayrhofer. A context authentication proxy for IPsec using spatial reference, December 2006. 52

- [MGH06] R. Mayrhofer, H. Gellersen, and M. Hazas. Security by spatial reference: Using relative positioning to authenticate devices for spontaneous interaction, 2006. submitted for publication. 41
- [Mic06] Microsoft. Genuine microsoft software f.a.q. list. <http://www.microsoft.com/genuine/downloads/faq.aspx>, 2006. Link checked 2006-08-23. 41
- [Mil06] Mark Samuel Miller. *Robust Composition: Towards a Unified Approach to Access Control and Concurrency Control*. PhD thesis, Johns Hopkins University, Baltimore, Maryland, USA, May 2006. 7
- [MNH<sup>+</sup>01] M. Maki, Y. Nishikawa, S. Hamada, M. Tokuda, and Y. Shimosio. Home information wiring system using a balanced cable for realizing a high-speed home-network. In *International Conference on Consumer Electronics, 2001.*, pages 236–237, Los Angeles, CA, USA, 19.–21. June 2001. ICCE. ISBN: 0-7803-6622-0, DOI: 10.1109/ICCE.2001.935290. 12
- [Mor02] Neal Morse. Wind at my back. Snow, August 2002. <http://www.spocksbeard.com/discography/snow.html#Anchor-15-3690>. xii
- [MS02] Kevin D Mitnick and William L Simon. *The Art of Deception: Controlling the Human Element of Security*. Wiley publishing, Indianapolis, 2002. [http://madchat.org/esprit/textes/The\\_Art\\_of\\_Deception.pdf](http://madchat.org/esprit/textes/The_Art_of_Deception.pdf). 35
- [Nie93] Jakob Nielsen. *Usability Engineering*. Morgan Kaufmann, 1993. 29, 46
- [NL94] Jakob Nielsen and Jonathan Levy. Measuring usability: preference vs. performance. *Commun. ACM*, 37(4):66–75, 1994. 21
- [Nor98] Donald Norman. *The Invisible Computer*. Number ISBN 0-262-64041-4. MIT Press, Cambridge, MA, pb edition, 1998. <http://www.jnd.org/books.html#434>. 6, 12, 28
- [Nor07] Donald Norman. *The design of future things (draft)*. Basic Books, New York, chapter 1 (draft 2007-mar-06) edition, 2007. <http://www.jnd.org/books.html#641>. 71
- [NS04] Tatsuo Nakajima and Ichiro Satoh. Personal home server: Enabling personalized and seamless ubiquitous computing environments. In *PERCOM '04: Proceedings of the Second IEEE*

*International Conference on Pervasive Computing and Communications (PerCom'04)*, page 341, Washington, DC, USA, 2004. IEEE Computer Society. 51

- [NSA02] Tatsuo Nakajima, Ichiro Satoh, and Hiroyuki Aizu. A virtual overlay network for integrating home appliances. In *Proceedings of the 2002 Symposium on Applications and the Internet (SAINT 2002)*, pages 246–253, Nara, Japan, 2002. ISBN: 0-7695-1447-2. INSPEC Accession Number: 7212499. Digital Object Identifier: 10.1109/SAINT.2002.994487. 14
- [NTH<sup>+</sup>06] Jun Noda, Mie Takahashi, Itaru Hosomi, Hisashi Mouri, Yoshiaki Takata, and Hiroyuki Seki. Integrating presence inference into trust management for ubiquitous systems. In *SACMAT '06: Proceedings of the eleventh ACM symposium on Access control models and technologies*, pages 59–68, New York, NY, USA, 2006. ACM Press. 51
- [NYHR05] Clifford Neuman, Tom Yu, Sam Hartman, and Ken Raeburn. The kerberos network authentication service (v5). RFC4120, July 2005. <http://tools.ietf.org/html/4120>, obsoletes RFC1510. 36
- [RHB03] Birgitte Ringbauer, Dr. Frank Heidmann, and Jakob Beisterfeldt. When a house controls its master—universal design for smart living environments. In *Proceedings of the Universal Access in HCI: Inclusive Design*, Crete, Greece, 2003. HCI. [hci.iao.fraunhofer.de](http://hci.iao.fraunhofer.de). 17, 18, 19
- [Rob91] Robert Bosch GmbH, Postfach 30 02 04, D-70442 Stuttgart, Deutschland. *CAN Specification 2.0B*, 2.0b edition, September 1991. <http://www.semiconductors.bosch.de/pdf/can2spec.pdf>. 8
- [Rok73] Milton Rokeach. *The nature of human values*. Free Press, 1973. New York. 17
- [SA99] Frank Stajano and Ross Anderson. The resurrecting duckling: Security issues for ad-hoc wireless networks. In *Security Protocols, 7th International Workshop Proceedings*, pages 172–194, 1999. <http://www.cl.cam.ac.uk/~fms27/papers/1999-StajanoAnd-duckling.pdf>. 34, 49, 52, 65, 73
- [SCFY96] Ravi S. Sandhu, Edward J. Coyne, Hal L. Feinstein, and Charles E. Youman. Role-based access control models. *IEEE Computer*, 29(2):38–47, Feb 1996. <http://csrc.nist.gov/rbac/sandhu96.pdf>. 37

- [Sch04a] Ben Schneiderman. Designing for fun: How can we design user interfaces to be more fun? *Interactions*, Sep-Oct 2004. <http://www.cs.umd.edu/~ben/Fun-p48-shneiderman.pdf>. 18
- [Sch04b] Barry Schwartz. *The Paradox of Choice: Why More Is Less*. Ecco, January 2004. 71
- [SG02] D. K. Smetters and Rebecca E. Grinter. Moving from the design of usable security technologies to the design of useful secure applications. In *NSPW '02: Proceedings of the 2002 workshop on New security paradigms*, pages 82–89, New York, NY, USA, 2002. ACM Press. <http://doi.acm.org/10.1145/844102.844117>. 15, 38, 45, 46
- [SGTGI04] Scarlet Schwiderski-Grosche, Allan Tomlinson, Swee Keow Goo, and James M. Irvine. Security challenges in the personal distributed environment. In *Vehicular Technology Conference, 2004. VTC 2004-Fall. 2004 IEEE 60th*, volume 5, pages 3267–3270, <http://personal.rhul.ac.uk/umai/274/Publications/Securitychallengesin2004.pdf>, September 2004. IEEE. 49
- [Sie07] Kathy Sierra. What comes after usability? Creating passionate users, January 2007. [http://headrush.typepad.com/creating\\_passionate\\_users/2007/01/what\\_comes\\_afte.html](http://headrush.typepad.com/creating_passionate_users/2007/01/what_comes_afte.html). 18
- [SJF<sup>+</sup>03] J. Seigneur, C. Jensen, S. Farrell, E. Gray, and Y. Chen. Towards security auto-configuration for smart appliances, 2003. 46
- [SS75] Jerome H. Saltzer and Michael D. Schroeder. The protection of information in computer systems (revised). In *Proceedings of the IEEE 63*, volume 9, pages 1278–1308, Cambridge, Mass. 02139, September 1975. Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, IEEE. <http://web.mit.edu/Saltzer/www/publications/protection/index.html>. 43
- [Sta00] Frank Stajano. The resurrecting duckling – what next? In *Lecture Notes in Computer Science*, volume 2133 / 2001, page 204, Cambridge, UK, April 2000. Security Protocols : 8th International Workshops, Springer Berlin / Heidelberg. ISSN: 0302-9743 <http://www.cl.cam.ac.uk/~fms27/papers/2001-Stajano-duckling.pdf>. 49, 53
- [Sta02a] Frank Stajano. *Security for Ubiquitous Computing*. John Wiley and Sons, February 2002. 34, 39, 76

- [Sta02b] Frank Stajano. Security for whom? the shifting security assumptions of pervasive computing. In Mitsuhiro Okada, Benjamin C. Pierce, Andre Scedrov, Hideyuki Tokuda, and Akinori Yonezawa, editors, *ISSS*, volume 2609 of *Lecture Notes in Computer Science*, pages 16–27. Springer, 2002. 2, 16, 39, 40, 69
- [SV04] Harm Anne Schotanus and Cor A. A. Verkoelen. Extended home environment from a security perspective. Technical report, The Netherlands Organisation for applied scientific research, Physics and electronics laboratory, The Hague, March 2004. <https://doc.telin.nl/dscgi/ds.py/Get/File-32077/>. 39, 41, 48, 49, 50, 52, 74
- [Tie] Joka kodin tietoturvaopas. <http://www.tietoturvaopas.fi/fi/dokkarit/JokaKodinTietoturvaopas.pdf>. 47
- [UNO03] Basic security in the field – staff safety, health, and welfare. CD-ROM and the web, <http://www.unops.org/security>, 2003. Complementary e-learning to UN Security in the field’ manual <http://www.reliefweb.int/library/documents/security.pdf>. 65
- [UPn03] *UPnP<sup>TM</sup> Device Architecture*, v1.0.1 draft edition, 2 December 2003. <http://www.upnp.org/resources/documents/CleanUPnPDA101-20031202s.pdf>. 21
- [VSSM01] Alladi Venkatesh, Norman Stolzoff, Eric Shih, and Sanjoy Mazumdar. The home of the future: An ethnographic study of new information technologies in the home. *Advances in Consumer Research*, Volume XXVIII:88–96, 2001. [crito.uci.edu, citeseer.ist.psu.edu/596768.html](http://crito.uci.edu/citeseer.ist.psu.edu/596768.html). 1
- [War05] Dr. Kevin Warwick. Project cyborg 1.0. Website of Kevin Warwick, 2005. <http://www.kevinwarwick.com/Cyborg1.htm>. 41, 51
- [Wei99] Mark Weiser. The computer for the 21st century. *SIGMOBILE Mob. Comput. Commun. Rev.*, 3(3):3–11, 1999. 11, 16, 20
- [Whi04] Alma Whitten. *Making Security Usable*. PhD thesis, School of Computer Science, Carnegie Mellon University, May 2004. 45
- [Wib03] Charlotte Wiberg. *A Measure of Fun*. PhD thesis, Umeå University, Sweden, 2003. <http://www.informatik.umu.se/~colsson/AvhandlingsCD/index.htm>. 18
- [Wib05] Charlotte Wiberg. Fun in the home: Guidelines for evaluating interactive entertainment on the web. In *In proceedings*

- of HCI International*, 12th international conference on Human Computer Interaction, Las Vegas, USA, July 2005. HCI International. <http://www.informatik.umu.se/~colsson/Articles/finalhciint2005CW.pdf>. 18
- [WSRE03] Stephen A. Weis, Sanjay E. Sarma, Ronald E. Rivest, and Donald W. Engels. Security and privacy aspects of low-cost radio frequency identification systems. In *International Conference on Security in Pervasive Computing, LNCS*, 2003. <http://saweis.net/pdfs/spc-rfid.pdf>. 38
- [WT98] Alma Whitten and Doug Tygar. Usability of security: A case study. Technical Report CMU-CS-98-155, Carnegie Mellon University, Pittsburgh, PA 15213, December 1998. [http://www.tygar.net/papers/Why\\_Johnny\\_Cant\\_Encrypt/Usability\\_case\\_study.pdf](http://www.tygar.net/papers/Why_Johnny_Cant_Encrypt/Usability_case_study.pdf). 44
- [WWB<sup>+</sup>05] Susan Wiedenbeck, Jim Waters, Jean-Camille Birget, Alex Brodskiy, and Nasir Memon. Authentication using graphical passwords: effects of tolerance and image choice. In *SOUPS '05: Proceedings of the 2005 symposium on Usable privacy and security*, pages 1–12, New York, NY, USA, 2005. ACM Press. 50
- [XAC03] XACML. A brief introduction to xacml. [http://www.oasis-open.org/committees/download.php/2713/Brief\\_Introduction\\_to\\_XACML.html](http://www.oasis-open.org/committees/download.php/2713/Brief_Introduction_to_XACML.html), 14 March 2003. 56
- [XRM01] Content Guard. *The need for a Rights Language*, technical white paper, version 1.0 edition, 2001. <http://www.xrml.org/reference/TheNeedForARightsLanguage.pdf>. 56
- [Yee04] Ka-Ping Yee. Aligning security and usability. *Security and Privacy Magazine, IEEE*, 2(5):48–55, Sept.-Oct. 2004. ISSN: 1540-7993. DOI: 10.1109/MSP.2004.64. California Univ., Berkeley, CA, USA. 27, 44, 45
- [ZK02] Kan Zhang and Tim Kindberg. An authorization infrastructure for nomadic computing. In *SACMAT*, pages 107–113, 2002. 8
- [ZS96] Mary Ellen Zurko and Richard T. Simon. User-centered security. In *NSPW '96: Proceedings of the 1996 workshop on New security paradigms*, pages 27–33, New York, NY, USA, 1996. ACM Press. 1, 37