



Designing Privacy-Preserving Personalized Public Display Systems

Morin Ostkamp

Geoinformatik

Designing Privacy-Preserving Personalized Public Display Systems

Inaugural-Dissertation
zur Erlangung des Doktorgrades der Naturwissenschaften
im Fachbereich Geowissenschaften
der Mathematisch-Naturwissenschaftlichen Fakultät
der Westfälischen Wilhelms-Universität Münster

vorgelegt von
Morin Ostkamp
aus Münster
— 2015 —

Dekan:	Prof. Dr. Hans Kerp
Erstgutachter:	Prof. Dr.-Ing. Christian Kray
Zweitgutachter:	Prof. Francis Harvey, PhD
Tag der mündlichen Prüfung:	27.11.2015
Tag der Promotion:	27.11.2015

Acknowledgements

Although there is just one name printed on the cover of this thesis, I could never have done this all by myself. Research is a collaborative process, so I am deeply grateful to everyone, who supported me in my work. First of all, I would like to thank my supervisors for their outstanding endeavors. Christian Kray, Francis Harvey, and Gernot Bauer helped me with their invaluable expertise to structure my thoughts, focus on key points, and keep me moving ahead.

During the last five years, I had the luck to work with nice people. I would thus like to thank all colleagues, co-authors, and especially—in alphabetical order—Matthias Böhmer, Tobias Brüggentisch, Ioannis Delikostidis, Thore Fechner, Holger Fritze, Sven Gehring, Peter Grobara, Achim Hennecke, Arne Kaiser, Silvio Kühn, Stefan Lösing, Sven Luzar, Champika Manel, Nicholas Schiestel, Matthias Seuter, Sonja Wälzlein, and Mirko Werner. Stefan Adam kindly allowed me to use some of his beautiful photos; the background of the cover is a dazzling example.

Most certainly, writing a doctoral thesis is a long-term project with ups and downs. I have been warned that there would be times when the going gets rough. Nevertheless, I was close to giving up at least once. This thesis proves, however, that I did not surrender in the end. The remarkable encouragement of priceless friends is what helped me to “keep calm and carry on.” Finally, I would like to thank my family for their unlimited and loving support.

Abstract

Digital public displays are a popular means of communication nowadays. They are commonly used as information outlets at traffic hubs, shopping malls, or public places in general. They feature some key advantages in comparison to other media types. For example, they are more flexible and more up-to-date than static paper-based approaches. Since they are always *situated* in a certain context, they can disseminate information tailored to a specific location. Compared to personal devices, that usually also have significantly smaller screens, digital public displays allow to distribute information within a spatial frame of reference. Moreover, they are visually prominent and provide broad accessibility.

Clearly, showing users content that is *relevant* to them is an important issue. For example, due to a lack of relevant content, many people have developed a blindness towards public displays. Personal content is often regarded as relevant, but that calls for certain means of privacy in turn. This thesis focuses on designing *privacy-preserving personalized public display systems*. It addresses three research questions: (1) What are main privacy threats on public displays? (2) What are countermeasures to those privacy threats? (3) How to support the design process of public displays?

Three tangible contributions address each research question: (1) The *STRIDED** *privacy threat model* for public displays is based on the renowned STRIDE model by Microsoft and the OWASP application se-

curity risks. The threat model helps to identify and prioritize privacy issues that personalized public display systems may be subject to. Along with the threat model comes a design space for privacy demands on public displays. (2) A *list and classification of existing countermeasures*, which is based on an extensive literature survey. These countermeasures can be used to address—at least some of—the privacy issues identified by the privacy threat model. Additionally, three novel countermeasures add to the list of existing measures. (3) A *novel approach to engineer public display systems* based on a realistic audiovisual simulation and a state-transition graph. The approach has been integrated in a holistic process, that provides a new methodology to design, prototype, and evaluate privacy-preserving personalized public display systems. This includes a systematic analysis of existing concepts to engineer public display systems, a novel approach that integrates many of the benefits of previous concepts, and an architecture for a toolkit implementing the approach.

Designers and researchers can use these contributions to create public displays, that do not pose a threat to the user's privacy: The threat model and the design space can be used to build privacy-aware public display systems that align with users' privacy perceptions and needs more closely. The list of countermeasures—which is also represented as a heat map—allows designers and researchers of public display systems to quickly identify the most commonly used countermeasure for a particular privacy threat. Finally, the integrated process can be applied directly, since it has been realized and published as a publicly available software toolkit. In conclusion, this thesis can thus contribute towards simplifying and accelerating the development of privacy-preserving personalized public display systems.

Zusammenfassung

Öffentliche Bildschirme — Digital-Public-Displays — sind heutzutage ein allgegenwärtiges Kommunikationsmedium. Sie werden oft als Informationsquellen an Verkehrsknoten oder in Einkaufshäusern sowie öffentlichen Plätzen im Allgemeinen genutzt. Verglichen mit anderen Medien bieten sie viele Vorteile. Sie sind z.B. flexibler und aktueller als statische Druckerzeugnisse. Da sie stets in einem örtlichen Kontext *eingebettet* sind, können sie für diesen Ort maßgeschneiderte Informationen veröffentlichen. Verglichen mit persönlichen Geräten, die meist nur über kleinere Bildschirme verfügen, können Digital-Public-Displays Informationen innerhalb eines räumlichen Bezugsrahmens verteilen. Zudem sind sie optisch auffallend und leicht zugänglich.

Benutzern *relevante* Inhalte zu präsentieren ist zweifelsohne wichtig. Da bisher aber relevante Inhalte oft fehlten, entwickelten Benutzer eine Blindheit gegenüber öffentlichen Bildschirmen. Persönliche Inhalte sind meist relevant, erfordern jedoch besondere Datenschutzmaßnahmen. Diese Arbeit konzentriert sich auf den *Entwurf personalisierter öffentlicher Bildschirme, die die Privatsphäre schützen*. Sie untersucht dazu drei Forschungsfragen: (1) Was sind die größten Gefahren für die Privatsphäre auf öffentlichen Bildschirmen? (2) Welche Gegenmaßnahmen existieren für diese Gefahren? (3) Wie kann der Entwurf öffentlicher Bildschirme unterstützt werden?

Drei konkrete Beiträge widmen sich je einer Forschungsfrage: (1) Das *STRIDED*-Modell für Gefahren* auf öffentlichen Bildschirmen basiert auf dem renommierten STRIDE-Modell von Microsoft und den OWASP-Application-Security-Risks. Das Gefahren-Modell hilft dabei, Risiken

für die Privatsphäre auf personalisierten öffentlichen Bildschirmen zu identifizieren und zu priorisieren. Gemeinsam mit dem Gefahren-Modell wird auch ein Design-Space für die Anforderungen an die Privatsphäre auf öffentlichen Bildschirmen vorgestellt. (2) Eine *Liste mit- samt einer Klassifizierung existierender Gegenmaßnahmen*, die auf einer umfassenden Literatur-Recherche basiert. Diese Gegenmaßnahmen adressieren — zumindest einige der — Risiken für die Privatsphäre, die das Gefahren-Modell identifiziert. Außerdem wird die Liste der vorhandenen Gegenmaßnahmen um drei neue Verfahren erweitert. (3) Ein *neuer Ansatz für die Entwicklung öffentlicher Bildschirme*, der auf einer realistischen audio-visuellen Simulation und einem Zustands-übergang-Graphen basiert. Der Ansatz wurde in einen ganzheitlichen Prozess integriert, der eine neue Methodik für den Entwurf, die prototypische Umsetzung und die Evaluation von öffentlichen Bildschirmen darstellt. Diese Methodik umfasst eine systematische Analyse vorhandener Konzepte für die Entwicklung solcher Bildschirme, einen neuen Ansatz, der viele Vorteile vorhergehender Konzepte vereint, und die Architektur für ein Toolkit, das diesen Ansatz umsetzt.

Designer und Forscher können diese Ergebnisse nutzen, um Systeme zu erstellen, die die Privatsphäre der Benutzer schützen: Das Gefahren-Modell und der Design-Space erlauben es, öffentliche Bildschirme zu entwerfen, die sich an der Auffassung von Privatsphäre und den Bedürfnissen der Benutzer orientieren. Die Liste der Gegenmaßnahmen — die zudem auch in Form einer Heat-Map dargestellt wird — erlaubt es Designern und Forschern, die für ein bestimmtes Risiko am häufigsten genutzte Gegenmaßnahme schnell zu identifizieren. Außerdem kann der integrierte Prozess direkt angewandt werden, da er als öffentlich verfügbare Software umgesetzt wurde. Zusammenfassend kann diese Arbeit dazu beitragen, die Entwicklung personalisierter öffentlicher Bildschirme, die die Privatsphäre schützen, zu vereinfachen und zu beschleunigen.

Nil desperandum.

Contents

Acknowledgements	i
Abstract	iii
List of Publications	xv
List of Figures	xvii
List of Tables	xxi
Remarks	xxv
I. Introduction	1
1. Overview	3
2. The Evolution of Public Displays	13
2.1. Ubiquitous Proliferation	14
2.2. Display Blindness	16
2.2.1. Operator’s Point of View	16
2.2.2. User’s Point View	18
3. Motivation	21
3.1. Addressing Display Blindness with Personalization . . .	22
3.2. Preserving Privacy During Personalization	25

4. Objectives	29
4.1. Research Questions	30
4.2. Scientific Contributions	31
4.3. Scope	34
4.4. Thesis Outline	35
II. Methodology, Key Concepts, and Related Work	37
5. Research Methodology	39
6. Key Concepts in Public Display Systems	43
6.1. Context	44
6.2. Privacy	47
6.3. Personalization	53
6.4. Design	54
6.5. Threats, Threat Models, and Countermeasures	57
7. Related Work on Public Display Systems	65
7.1. Context	66
7.2. Privacy	68
7.3. Personalization	77
7.4. Design	79
7.5. Threats, Threat Models, and Countermeasures	89
7.5.1. Literature Survey	90
7.5.2. Visualization of the Surveyed Work	92
7.5.3. Outcomes and Discussion	97
7.5.4. Limitations	118
7.5.5. Summary	119

7.6.	Toolkits and Frameworks	120
7.6.1.	Design	123
7.6.2.	Content	126
7.6.3.	Interaction	129
7.6.4.	Social Connections	132
8.	Summary	137
III.	Designing Privacy-Preserving Personalized PDS	141
9.	Challenges	143
9.1.	Situatedness	144
9.2.	Form Factors	146
9.3.	Fixed Environmental Factors	148
9.4.	Dynamic Environmental Factors	150
9.5.	Mobile Devices	152
9.6.	Multi-Display Networks	154
9.7.	Acceptance	156
9.8.	Legal Constraints	158
10.	Approaches	161
10.1.	Privacy Threat Model	162
10.1.1.	Deriving a Theoretical Grounding	162
10.1.2.	Derived Design	181
10.2.	Countermeasures	203
10.2.1.	Visual Multiplexing	205
10.2.2.	Visual Highlighting	218
10.2.3.	Visual Interaction	233

10.3. Process Integration	242
10.3.1. Immersive Public Display Evaluation and Design Toolkit	243
10.3.2. Immersive Video Environment	256
11. Prototypes	261
11.1. Privacy Threat Model	261
11.2. Countermeasures	265
11.2.1. Visual Multiplexing	265
11.2.2. Visual Highlighting	278
11.2.3. Visual Interaction	282
11.3. Process Integration	285
11.3.1. Immersive Public Display Evaluation and Design Toolkit	285
11.3.2. Immersive Video Environment	303
12. Evaluation	309
12.1. Privacy Threat Model	310
12.2. Countermeasures	321
12.2.1. Visual Multiplexing	321
12.2.2. Visual Highlighting	331
12.2.3. Visual Interaction	346
12.3. Process Integration	350
12.3.1. Immersive Public Display Evaluation and Design Toolkit	350
12.3.2. Immersive Video Environment	354
13. Summary	359

IV. Reflections	361
14. Discussion	363
14.1. Privacy Threat Model	364
14.2. Countermeasures	367
14.2.1. Visual Multiplexing	367
14.2.2. Visual Highlighting	375
14.2.3. Visual Interaction	381
14.3. Process Integration	386
14.3.1. Immersive Public Display Evaluation and Design Toolkit	387
14.3.2. Immersive Video Environment	392
14.4. Summary	395
15. Conclusion	401
15.1. Contributions	401
15.2. Future Work	408
V. Appendix	411
Bibliography	413
Supplementary Material	443
Student Theses	443
Privacy Threat Model Relations	444
Cover Letter for the Qualitative Evaluation of CI	451
Curriculum Vitae	453

List of Publications

The major part of the work reported on in this thesis has been published in these peer-reviewed papers:

[175] Ostkamp, M., Kray, C., Bauer, G. Towards a Privacy Threat Model for Public Displays. Proc. EICS '15.

[171] Ostkamp, M., Heitmann, S., Kray, C. Short-range optical interaction between smartphones and public displays. Proc. PerDis '15.

[239] Wilhelm, D., Fechner, T., Ostkamp, M., Kray, C. Natural interaction with video environments using gestures and a mirror-image avatar. Proc. Interact '15.

[174] Ostkamp, M., Kray, C. Supporting Design, Prototyping, and Evaluation of Public Display Systems. Proc. EICS '14.

[176] Ostkamp, M., Luzar, S., Bauer, G. QR Codes on Curved Media Facades — Two Approaches For Inverse Distortion Based on Raytracing and Image Warping. Proc. GRAPP '14.

[172] Ostkamp, M., Hülsermann, J., Kray, C., and Bauer, G. Using mobile devices to enable visual multiplexing on public displays: Three approaches compared. Proc. MUM '13.

[173] Ostkamp, M., and Kray, C. Prototyping mobile AR in immersive video environments. Workshop on Designing Mobile Augmented Reality, MobileHCI '13.

[170] Ostkamp, M., Bauer, G., and Kray, C. Visual Highlighting on Public Displays. Proc. PerDis '12.

[30] Böhmer, M., Gehring, S., Löchtefeld, M., Ostkamp, M., and Bauer, G. The Mighty Un-touchables — Creating Playful Engagement on Media Facades. Proc. MobileHCI '11.

[169] Ostkamp, M. and Bauer, G. Multipleye — Concurrent Information Delivery on Public Displays. Adjunct Proc. EuroITV '11.

List of Figures

1.1.	Examples of Public Displays	4
1.2.	Public Displays at Times Square	7
4.1.	Thesis Structure	36
6.1.	The Spiral Model by Boehm	55
7.1.	The Audience Funnel and the Honeypot Effect	84
7.2.	Literature Survey: Hive Plot	95
7.3.	Literature Survey: Applications and Threats	96
7.4.	The P-LAYERS Framework	124
10.1.	OWASP: Application Security Risks	163
10.2.	STRIDE Study: Relevances	175
10.3.	Privacy Threat Model: Final Design	181
10.4.	Privacy Threat Model: Visualization	198
10.4.	Privacy Threat Model: Visualization (continued)	199
10.4.	Privacy Threat Model: Visualization (continued)	200
10.4.	Privacy Threat Model: Visualization (continued)	201
10.5.	SIGCHI Topics of Interest	204
10.6.	Visual Multiplexing: FDM Example	213
10.7.	Visual Multiplexing: CDM Example (First Version)	216
10.8.	Visual Multiplexing: CDM Example (Second Version)	216
10.9.	Visual Multiplexing: TDM Example	218

10.10.	Visual Highlighting: Cardboard Analogy	220
10.11.	Visual Highlighting: Design	226
10.12.	Visual Interaction: Design	241
10.13.	IPED Toolkit: Workflow	254
11.1.	IPDPTM: Screenshot	263
11.2.	Visual Multiplexing: FDM Web App	273
11.3.	Visual Multiplexing: CDM Web App	274
11.4.	Visual Multiplexing: TDM Web App	275
11.5.	Visual Multiplexing: FDM Mobile App	276
11.6.	Visual Multiplexing: CDM Mobile App	276
11.7.	Visual Multiplexing: TDM Mobile App	277
11.8.	Visual Highlighting: Web App	280
11.9.	Visual Highlighting: Mobile App	281
11.10.	Visual Interaction: Image Processing Chain	283
11.11.	IPED Toolkit: Architecture	286
11.12.	IVE: Panoramic Video Footage	288
11.13.	IPED Toolkit: Backend Overview	292
11.14.	IPED Toolkit: Creating Routes	293
11.15.	IPED Toolkit: Creating Overlays	294
11.16.	IPED Toolkit: Third Party Content	297
11.17.	IPED Toolkit: Remote Control	299
11.18.	IVE: Mirror Image Avatar	305
11.19.	IVE: Avatar Control Gestures	306
12.1.	IPDPTM: Students' Models	317
12.1.	IPDPTM: Students' Models (continued)	318
12.1.	IPDPTM: Students' Models (continued)	319
12.1.	IPDPTM: Students' Models (continued)	320
12.2.	Visual Multiplexing: Studied Content Types	322
12.3.	Visual Multiplexing: Answers by Content Type	329

12.4.	Visual Multiplexing: Answers by Method	329
12.5.	Visual Multiplexing: NASA TLX by Content Type	330
12.6.	Visual Multiplexing: NASA TLX by Method	330
12.7.	Visual Highlighting: Spider Diagram	335
12.8.	Visual Highlighting: Icon Grids	337
12.9.	Visual Interaction: Signals by Distance	349
12.10.	IPED Toolkit: Students' Screenshots	353
14.1.	Visual Multiplexing: Participant Using TDM	372

List of Tables

5.1.	Applied Study Types	40
6.1.	Taxonomy of Privacy	50
6.2.	Research Paradigms and Questions	56
7.1.	Personalization Usage Models	78
7.2.	Public Display User Values	86
7.2.	Public Display User Values (continued)	87
7.2.	Public Display User Values (continued)	88
7.3.	Literature Survey: Overview	93
7.3.	Literature Survey: Overview (continued)	94
7.4.	Literature Survey: Application Scenarios	99
7.5.	Literature Survey: Matches	102
7.6.	Literature Survey: Mismatches	102
7.7.	Literature Survey: Personalization Usage Models	103
7.8.	Literature Survey: Privacy Threats	105
7.9.	Literature Survey: Countermeasures	109
7.10.	Literature Survey: Heat Map	110
7.11.	Literature Survey: Research Opportunities	114
10.1.	STRIDE Study: Hypotheses	165
10.2.	STRIDE Study: Public Display Actions	167
10.3.	STRIDE Study: Structure	169
10.3.	STRIDE Study: Structure (continued)	170

10.3.	STRIDE Study: Structure (continued)	171
10.4.	STRIDE Study: Total Likert Scores (Questions)	172
10.5.	STRIDE Study: Rater's Rankings	173
10.6.	STRIDE Study: Total Likert Scores (Applications)	174
10.7.	OWASP: Threat Agent Factors	183
10.8.	OWASP: Likelihood and Impact Levels	184
10.9.	Privacy Threat Model: Agents	184
10.10.	Privacy Threat Model: Threats	185
10.10.	Privacy Threat Model: Threats (continued)	186
10.10.	Privacy Threat Model: Threats (continued)	187
10.10.	Privacy Threat Model: Threats (continued)	188
10.10.	Privacy Threat Model: Threats (continued)	189
10.11.	Privacy Threat Model: Weaknesses	192
10.12.	Privacy Threat Model: Countermeasures	193
10.12.	Privacy Threat Model: Countermeasures (continued)	194
10.12.	Privacy Threat Model: Countermeasures (continued)	195
10.12.	Privacy Threat Model: Countermeasures (continued)	196
10.13.	Privacy Threat Model: Effects	197
10.14.	Visual Highlighting: Comparison Criteria	231
10.14.	Visual Highlighting: Comparison Criteria (continued)	232
10.15.	Evaluation and Design Methods	247
10.15.	Evaluation and Design Methods (continued)	248
11.1.	Visual Interaction: QR Code	283
12.1.	IPDPTM: NASA TLX	321
12.2.	Visual Multiplexing: Questions and Hypotheses	323
12.3.	Visual Highlighting: Quality	332
12.4.	Visual Highlighting: Quantity	333
12.5.	Visual Highlighting: Robustness	334
12.6.	Visual Highlighting: Questions and Hypotheses	338

12.7.	Visual Highlighting: Structure	340
12.8.	Visual Highlighting: Efficiency	344
12.9.	Visual Highlighting: Effectiveness	344
12.10.	Visual Highlighting: NASA TLX	345
12.11.	Visual Interaction: System Characteristics	346
12.12.	Visual Interaction: Signals by Smartphone	348
12.13.	Visual Interaction: Multi-User Analysis	348
12.14.	IPED Toolkit: NASA TLX	352
12.15.	IPED Toolkit: UMUX	354
14.1.	Visual Highlighting: Benefits and Drawbacks	380
15.1.	Scientific Contributions	406
15.2.	Practical Contributions	407

Remarks

According to the recommendations of the “Publication Manual of the American Psychology Association,” this thesis includes citations embedded within the original material [13, p. 173]. These embedded citations may have the same appearance as the genuine citations within the body of this thesis. However, they should not be confused, as the numbers of the embedded citations may not exist or match the actual references in the appendix.

Unless otherwise specified, all figures are property of Morin Ostkamp.

I

Introduction

1 Overview

The term “public display” describes a plethora of things. Verbalizing it may help to recognize this: to display something publicly. Thus, a public display may refer to a signpost, a stake, an artwork, or any other object exposed in a public place. The term may also relate to an action, like the public display of love or affection expressed by a kiss in public. In this thesis, however, a *public display system (PDS)* is a digital appliance, most often a computer monitor of arbitrary size, showing content of any kind. Figure 1.1 shows a number of examples.

Public displays turned into ubiquitous companions in the course of the last years. Some readers, especially those who are actively engaged in current research on such displays, may be very familiar with statements like the last one. Many publications begin by stating this observation, and even though this introduction may appear trite and worn out, it is true though. Readers may put this thesis down for a moment and take a look around: In all likelihood there are public displays around. The following three examples taken from the author’s real everyday life may underpin this:

Getting to Work. It is 7:30 in the morning. Morin hastily leaves his home to catch his train. Quite exhausted, he arrives at the station, only to find a public display telling him that his train will



Figure 1.1.: Examples of public displays as addressed in this thesis.

be 10 minutes late. Later on in the train, another public display monitors the current speed, next stops, and connecting trains. When the train arrives at Münster Central Station, Morin gets out of the train, walks down the platform, and arrives at the entrance hall. A large public display casts some news and ads in his direction. On his way out, he passes a telephone booth equipped with a display that can be used to look up phone numbers or to send text messages. But Morin notices those screens only marginally, as he is in a hurry to catch his bus—in vain. A public display at the bus stop says that the next bus will arrive in 20 minutes. While waiting, Morin enjoys some artistic videos shown on public screens in the vaudeville across the street. When the bus arrives, Morin steps in and he directly peeks at the screen mounted on the ceiling. It shows the current time as well as the following stops, and Morin realizes that he will be late. Twenty minutes

later, Morin arrives at the institute for Geoinformatics. While rushing in, he does not even think about the two public displays he has just passed: One of them would have informed him about a scheduled power outage and the other one would have told him a joke to cheer up his mood.

Getting Food. Around 1:00 p.m. Morin is almost starving, as he worked really hard all morning to make up for the lost time. He leaves his office and heads to the refectory. Right in the middle of the foyer, there are four large public displays hanging from the ceiling. Two of them show an overview of the daily menu, while the remaining two advertise goods from the campus store and upcoming events. Once Morin made up his mind about what to eat today, he enters the dining hall. Public displays are mounted atop of each counter, showing the names and photos of the dishes to get there. Morin picks the counter with the most mouth-watering photo, only to find that the actual serving has little resemblance with that.

Shopping. After work, Morin decides to run some errands on his way home. Surprisingly, there are no public displays to welcome him at the entrance of the supermarket. Subconsciously, Morin notices the lack of eye-catching ads with delight. He makes his way to the cheese counter, preparing to wait in line. While waiting, his eyes wander from the counter to the storeroom with the large window pane and—eventually—to the public display attached to the top of that window pane. He did not notice it at first glance as its appearance is subtle. A static ad, looking like it was written on chalk board, tells him to take advantage of a special deal on Gouda. What a lucky coincidence, as this is Morin's favorite kind of cheese. Or maybe not a coincidence at all?

Hopefully, these three authentic examples—not inspired by, but taken from the author’s real life—illustrate how public displays have proliferated our urban lives. Especially the last example shows how to subtly blend in such displays to a specific situation.

The most profound technologies are those that disappear. They weave themselves into the fabric of everyday life until they are indistinguishable from it. —*Mark Weiser [237]*

This is what Weiser proclaims in his famous article [237] published in 1991. Apparently, however, Weiser’s vision has not completely become reality yet. Most public displays are not subtly integrated into daily routines to let them disappear. In contrast, many public displays struggle to catch our attention: Flamboyant colors, flashy animations, and prominent placings are just three examples of employed means. A probably well-known place hosting this type of public displays is the Times Square in New York City, see Figure 1.2. Hundreds of screens in various shapes and sizes compete for attention. This development is probably diametrically opposed to Weiser’s view of the future:

Most important, ubiquitous computers will help overcome the problem of information overload. There is more information available at our fingertips during a walk in the woods than in any computer system, yet people find a walk among trees relaxing and computers frustrating. Machines that fit the human environment instead of forcing humans to enter theirs will make using a computer as refreshing as taking a walk in the woods. —*Mark Weiser [237]*

Researchers found that most people nowadays tend to see a cause of *information overload* in public displays [154]. As a consequence of this display-induced information overload, people apparently developed an aversion towards such public displays. Based on these findings, Müller et al. coined the term *display blindness* [156].



Figure 1.2.: Public displays at the Times Square in New York City. Picture taken by Stefan Adam.

Nevertheless, *digital signage*, which is a particular form of public displays, mainly focused on application scenarios related to advertisement and product placement, continued to grow significantly. According to market reports and experts reviews [3] and Lyle Bunn¹, the market for digital signage is forecasted to reach an estimated compound annual growth rate (CAGR) of 8.94% from 2014 to 2020. Aside

¹In his book on digital signage, Jimmy Schaeffler calls Lyle Bunn “one of the better known and better versed champions of the recent digital signage movement” [195, p. 5].

from these figures being mostly relevant to decision makers, end users may notice this trend in their everyday lives, as more and more conventional signage will be replaced by its digital counterpart.

This shift from analog to digital is fostered by prices on the decline and fast-paced technological advances. While large, yet compact, computer screens based on LED technology, for example, become less expensive year by year, their visual fidelity is constantly increasing. The calm and steady visual appearance of flat screens (in contrast to the flickering images of old-fashioned cathode-ray monitors) probably also gave some impetus to this development. Optimized power consumption is another relevant topic addressed by current research: Electronic ink, for example, is an approach that allows to drive tablet devices, such as e-paper readers, for a long time. The same technology could be used to run less dynamic signage, as described in the shopping example scenario above, without generating excessive amounts of heat or consuming innumerable kilowatts.

This chapter contoured the ubiquitous proliferation of public displays and introduced the term display blindness—Chapter 2 explains both aspects in more depth. The overall objective of this thesis is to address the root of this blindness in order to mitigate its negative effects. Personal content is often regarded as a viable approach, but that calls for certain means of privacy in turn. The intended audience of this thesis can thus be broken down into the following four segments:

Researchers. Public displays have drawn the attention of numerous researchers and research groups around the world. There is a large body of research on personalized public displays, and an even larger one on public displays in general. This thesis provides an extensive review of related work in the context of *personalization* and *privacy* by proposing an original classification scheme. This classification scheme provides a new perspective on the domain

and may support novice researchers in collecting and applying related work. The thesis also contributes to the research community by proposing (i) an (interactive public display) *privacy threat model*, (ii) a number of measures to counter the *threats* of that model, and (iii) an integrated process for designing, prototyping, and evaluating *privacy-preserving* public display systems.

Designers. Public displays are basically monitors connected to computers, running a special software to drive the content. Just like any other software, it abides to a certain life cycle established by the discipline of software engineering. Three important cornerstones of this life cycle are design, implementation, and evaluation. During the development of a public display system, these three phases may be especially challenging. This thesis presents an integrated process for designing, prototyping, and evaluating privacy-preserving public displays, which may be used to simplify these tasks. This thesis also provides tangible prototypes that support designers of privacy-preserving public display systems in identifying and tackling possible *privacy threats*.

Privacy advocates. The world we live in is under constant change. Fast-paced advances in technologies such as wireless communication and shrinking *form factors* have led to a densely interconnected *web of things*. Constant surveillance, e.g., by tracking someone's mobile phone or car, is not merely a dystopian view of the future anymore². Moreover, privacy has become a controversially discussed topic in the general public initiated by the Snowden revelations in mid 2013. Privacy advocates may benefit from reading this thesis, as it puts a special focus on privacy with regard to personalized public displays. One of its contributions is an (interac-

²In the year 2015, major car manufacturers already offer services like "BMW Assist," or "Mercedes-Benz mbrace," that track the position of a car and place an emergency call in case of an accident.

tive public display privacy) threat model. This web-based application is publicly available and may be used to analyze possible privacy issues of various systems in a “structured and systematic way,” cf. Section 11.3. The proposed set of *countermeasures* may also be of interest to this audience.

Everyone. Admittedly, this may seem flashy at first glance. However, one particular argument that this thesis tried to put forward so far is, that public displays have become a prevalent and ubiquitous experience in urban life nowadays. The three examples presented above hopefully helped to make this point: The examples depict real scenarios, actually experienceable in the year 2015. Thus, as public displays keep permeating urban environments, everyone should be concerned about privacy issues with regard to personalized public display systems and privacy in general. This thesis may help readers to grasp the possibilities offered by current technology and anticipate future developments. It may help to understand the importance of privacy in our everyday life, that has been saturated with ubiquitous computing devices— public displays being only one specific type:

Our most familiar ways of managing privacy depend fundamentally on features of the spatial world and of the built environment, whether that be the inaudibility of conversation at a distance, or our inability to see through closed doors. We can also rely on others to honor behavioral norms around physical touch, eye contact, maintenance of interpersonal space, and so on [6]. With information technology, our ability to rely on these same physical, psychological and social mechanisms for regulating privacy is changed and often reduced.

—*Palen and Dourish [178]*

At the sidelines, this thesis strives to weaken a very common refutation against any privacy issues uttered by the concerned: “I have nothing to hide” [213].

[...] The problem with the nothing to hide argument is the underlying assumption that privacy is about hiding bad things. [...] [It] stems from a faulty “premise that privacy is about hiding a wrong.”⁷⁵ —*Daniel Solove [213]*

Thus, even readers who think the topic covered in this thesis is no concern of them might reconsider their point of view eventually.

However, some readers may be tempted to stop reading at this point, as their interests are in other topics covered by related work. For example, if looking for literature on public displays or digital signage in general, the comprehensive overview by Jimmy Schaeffler [195] could be a good start. Readers seeking for inspiration on new application scenarios could enjoy the work by Florian Alt [5] about pervasive advertising on public displays.

The work by Marko Jurmu [109] is concerned with interactive multi-purpose systems, while Simo Hosio focuses on social networking on public displays. If the reader’s emphasize is on approaches towards context-adaptive systems, the work of Jörg Müller [150] makes good reading. If the impact of public displays on the communities they are deployed in is of interest, see the work of Nemanja Memarovic [142].

The remainder of the first part is structured as follows: First, there is a look at the evolution of public displays and their concomitant ubiquitous proliferation. Next, the motives that drive this thesis are presented. Finally, the last section highlights the objectives, research questions, and contributions.

2

The Evolution of Public Displays

As explained in the previous chapter, the term public display may refer to a number of things. In the context of this thesis, however, it relates to digital computer screens or monitors. As with any technology, these systems are subject to a fast paced development. Recent advances in processing speed, memory capacity, graphics, and network infrastructure have had a significant impact. Video games, for example, looked way different ten years ago, so did videos on the Internet with regard to quality. Another important landmark in the history of public display probably is the advent of flat screens, e.g., Plasma, LCD, or LED monitors. In particular, this technology allowed for new form factors, and the decreasing production costs paved the way for a true permeation of such displays. At the time of writing this thesis, organic display technologies (OLED) foreshadow the possibilities of the next screen generation, e.g., curved high-resolution displays.

Due to these advances in technology introduced in the early years of the twenty-first century, this thesis focuses on public displays and corresponding scientific publications that appeared after the year 2000. Similarly, Ardito et al. [15] narrowed down their survey on *interaction* with large displays to the years 2000–2014 and only presented

aggregated data on systems published prior to this timeframe. Nevertheless, the history of public display reaches back to the early 1970s, as described by Jimmy Schaeffler [195, p. 41]. Readers interested in more details may enjoy Schaeffler's book or the comprehensive breakdown of the historical development presented by Davies et al. [62].

Clearly, public displays have become popular artifacts of our everyday life. The examples and pictures provided in the previous chapter tried to underpin this fact. To put this into numbers, Schaeffler estimated that there had been about half a million public displays in North America in the year 2008 [195, p. 43]. He also expected this number to double until 2011. Section 2.1 addresses this ubiquitous proliferation in more depth. This noticeable penetration may, however, have had impacts on society, as discussed in Section 2.2.

2.1. Ubiquitous Proliferation

Besides prominent places, such as the Times Square in New York City, see Figure 1.2, or Shibuya Crossing in Tokyo, public displays have also permeated other urban [164, 180] or rural areas [222]. The advent of flat and small screens propelled this trend, as public displays can nowadays be installed even if there is little space available, such as in subway train cabins or elevators.

This gain in technical flexibility was accompanied by falling hardware prices [195, p. 12]. This combination led to a noticeable increase in public displays or digital signages in many countries around the world. This increase was likely caused by operators following a “me-too strategy” to catch up with the latest technological developments introduced by competitors. In contrast to this evolution, however, the application scenarios and use cases for public displays did not evolve

in a comparable fashion. In their extensive review of research on pervasive displays, Davies et al. note that “the applications conceived for public displays have shown remarkable resilience to change. [...] It appears that despite radical change in many application areas, potential display users are still drawn to the same set of applications that were conceived over a decade ago” [62, pp. 92–93].

One particular use case for public displays has been without doubt the most common one ever since: showing advertisements, as many studies and observations confirm [5, 74, 98, 143, 164, 195, 230]. Using such displays for commercials may seem reasonable from an economic point of view. Compared to other conventional ways of publication, e.g., analogous posters or bill-boards, their digital equivalents can be remotely controlled, support other content types besides static images, and allow for interaction. Moreover, the content shown on public displays can be updated easily, with little delay, and at little cost since no physical items have to be exchanged. This characteristic led many operators of ticket vending machines or similar kiosk systems to repurpose their screens by showing advertisements in case of inactivity, e.g., when no one is buying a ticket. As a consequence, citizens were soon exposed to a plethora of digital advertisements, employing different visual and acoustic means to compete for attention. A dystopian exaggeration of the real situation in 2002 is depicted in the movie “Minority Report:” Personalized advertisements are spread throughout an entire city and pursue the main character at every turn. This is clearly not the reality of today, but the fictive character and real citizens seem to have something in common: display blindness.

2.2. Display Blindness

During the day, people are exposed to a flood of visual, acoustic, olfactory, gustatory, or haptic stimuli. As the capacity of the brain is limited, not every impression is processed, but some are sorted out. This filtering process is complex and may be influenced by many things or over a certain period of time—a detailed presentation is, however, beyond the scope of this thesis. Pervasive advertising on public displays has been around for a relatively short period of time, compared to other historic impacts. However, it may already have had an influence on how people perceive public displays as noted by Müller et al. [156]. In their study, they analyzed how people react to such displays and contents. Their results indicate, that people tend “to ignore public displays when they expect nothing interesting to be presented” [156]. This behavior seems to be related to the one observed in studies on advertisements on websites, which is described by the term *banner blindness* [41]. Similarly, Müller et al. coined the term *display blindness* to describe the phenomenon witnessed in their studies.

This raises the question about why people tend to expect the contents of public displays to be uninteresting or even boring [156]. This question can be looked at from two sides: the public display operator and the public display user. The remainder of this section will describe each point of view in more detail.

2.2.1. Operator’s Point of View

In the context of this description, the term operator represents multiple stakeholders as identified by Alt et al. [5, 7, 8]: *display owner*, *space owner*, *display provider*, and *content provider*. These stakeholders may

share a common interest in maximizing the number of users. The following examples may help to clarify this.

Display owners usually buy and operate public displays for a certain purpose, for example, advertising retail goods in their shops. Another use case would be a public office that strives to provide citizens with answers to commonly asked questions prior to seeing a civil servant in order to minimize waiting and processing times. These stakeholders miss their goal if users do not look at their public displays due to display blindness.

Sometimes, space owners and display owners may be the same person, e.g., a store owner. Sometimes, however, these two stakeholders may be separate persons. For example, innkeepers may rent public displays (from display owners) and install them in their business as a courtesy in order to enhance their customers' experience. The negative perception of such displays as described by Müller et al. [156] may in fact lead to the opposite effect.

Display providers are companies that sell public displays to other stakeholders, for example, display owners or space owners. Display providers assemble the hardware and may also provide the software that can be used to orchestrate the content presentation. Companies such as Ströer, Wall, or VIDERO are examples of display provider. Since their business is the sales of such systems, the phenomenon of display blindness may negatively impact their sales figures, as customers could become reluctant to invest in more or newer public displays.

Content providers are in charge of aggregating, editing, and scheduling the contents shown on public displays. Their main goal is to reach as many users as possible since their payment is usually based on the estimated amount of *viewers* (sometimes referred to as *coverage* or *media penetration*). Their business is comparable to the one of display

providers, with regard to the marketing of public displays in various *locations*, e.g., subway stations, doctor's offices, or retail stores. Similarly, their business can be negatively affected by display blindness.

In summary, each stakeholder introduced above may regard the phenomenon of display blindness as a negative business impact. Consequently, they may strive to find ways to alleviate this effect as noted by Elhart et al.: "In order to make such displays more attractive, both researchers and advertisers have recently begun to explore the concept of interactive applications that allow passers-by to directly or indirectly control a display's content" [74]. The same authors also remark, however, that harmonizing the interests of users and operators may be challenging, as both parties may have opposing requirements. The next section analyzes display blindness from the user's perspective.

2.2.2. User's Point View

A particular product, be it hardware or software, is usually designed and built to fulfill a specific purpose. Ideally, the individual purpose corresponds to the user of that product. In that case the user's requirements and the features of the actual product should match well. When talking about public displays, however, the people that actually interact with these systems may not be the users whose requirements were considered: The stakeholders introduced above may, for example, *use* public displays to pursue a certain goal, such as advertising goods in retail stores. To differentiate between the stakeholders of public displays and the general public as the actual users of public displays, Alt et al. refers to the latter group as viewers [5, 8], which implies a level of passiveness. With regard to personalized public display systems, however, this term may be less appropriate: As explained in the remainder of this thesis, see Sections 3.1, 6.3, and 7.3, personalization

always requires some form of interaction, which implies a certain degree of activity. Thus, literature—including this thesis—often refers to the general public as the users of such displays, as these people are exposed to the contents of those systems and may interact with them.

As mentioned above, research found that users tend to pay little attention to public displays [98]. In their paper on display blindness, Müller et al. [156] try to provide some psychological background on the observed lack of attention:

In other [research] areas, lack of attention for aspects of the environment has been explained by the fact that attention is highly selective. The world provides far too much information to be processed by an individual. This is especially true for urban environments, where Milgram [6] showed that many individuals experience information overload. Milgram identified six common reactions to information overload, among them the allocation of less time to each input and disregard of low-priority inputs. In their survey on information overload, Eppler and Mengis [2] define the concept as follows: “Information overload describes the situation when too much information affects a person and the person is unable to recognize, understand or handle this amount of information.” They conclude that when information supply exceeds information-processing capacity, a person has difficulties in identifying relevant information. He/she becomes highly selective and ignores large amounts of information, has difficulties in identifying relationships between details and overall perspective and needs more time to reach a decision. —Müller et al. [156]

In summary, users may regard the phenomenon of display blindness as a positive concept. It may help them to quickly differentiate between relevant and irrelevant information to efficiently work on a specific task at hand. Consequently, Alt et al. suggest “that interactivity has the potential to overcome this phenomenon [of display blindness]” [11]. In a similar vein, Davies et al. conclude that in order “to provide content [that is] relevant to passersby, displays must offer sophisticated personalization” [65]—one of the catalysts of this thesis. Relevant content may help users to solve problems and execute tasks efficiently, while rebuilding their expectations of and regaining their trust in public displays. Addressing the issue of display blindness is thus a challenge both parties, i.e., operators as well as users, should strive to take on.

3

Motivation

The previous section discussed the phenomenon of display blindness, which is rooted in the users' negative experiences with public displays or the shown contents, respectively. Yet, public displays do not have to be solely annoying experiences after all, as hopefully illustrated by the three examples presented in Chapter 1. Schmidt et al. consider public displays to be one of six current technologies that are most likely to become a ubiquitous experience within the next twenty years:

Whereas nowadays digital signage often shows mere adaptations of traditional content, networking capabilities as well as sensors will allow content to be easily updated and adapted to the audience, potentially making public displays a future communication medium.³ A key challenge is to create a pleasant and convenient user experience that fosters people's engagement with public displays.

—Schmidt et al. [199]

At first sight, the swift proliferation of personal mobile devices, such as smartphones, may have defeated the purpose of public displays. Information can nowadays be pushed to the user's device and can thus be routed in a very directed manner. Despite this development, public displays may still serve a valuable purpose when “addressing groups

of people that can be found in a certain location” [8]. Hamhoun and Kray analyzed the use of public displays in such situations, e.g., sport matches, festivals, fairs [89] and pilgrimages [90]. They point out that using smartphones in crowded areas may actually be hazardous and that the required technical infrastructure may not scale well enough to serve all users simultaneously. The results of their studies indicate, that dynamic signage systems may support users in these situations, alleviate the aforementioned problems, and actually increase the overall safety.

It may thus be worthwhile for researchers to focus on public displays, even though other ubiquitous technologies, such as the smartphone, emerged in recent years. In any case, it is crucial to concentrate on aspects of particular importance from the user’s point of view. Personalization appears to be an approach commonly regarded as encouraging. Section 3.1 discusses this approach in more depth. However, focusing on personalization alone may not suffice since privacy implications also appear to have a major impact on public display systems as discussed in Section 3.2.

3.1. Addressing Display Blindness with Personalization

Personalization of public displays is a complex issue; there exists a large body of research on this topic. Section 6.3 provides more details, for example, by introducing the three *personalization usage models* as proposed by Davies et al. [64]. This section, however, tries to motivate how personalization can help to approach the effect of display blindness in general. Though there are different approaches to personalizing public displays, e.g., user-generated content and emotional binding, researchers seem to agree on the overall appropriateness of the basic notion:

To provide content relevant to passersby, displays must offer sophisticated personalization. —*Davies et al. [65]*

Support for interaction and user-generated content represents a promising direction towards more valuable digital public displays. —*José et al. [108]*

[...] One of the key challenges with big displays is to provide interesting content and to design engaging activities. User-generated content in public spaces with carefully tailored interaction options could very well be the answer to this challenge. —*Hosio et al. [97]*

We can even say that people wanted to see personalized and situated content, i.e., content according to their preferences. [...] This could be provided through a menu where people could choose their preferred category [...]. A more advanced approach would be to use user profiles to select the information automatically. —*Memarovic et al. [143]*

The most noticeable difference between the approaches may be the content type. There are at least two manifestations: private content (such as, e.g., messages [38, 101], calendars [49, 231], or directions [119, 167]) and user-generated content (such as, e.g., multi-media assets [7, 8, 108, 224], adaptive profiles [6, 152], or community-related data [9, 48, 157]). In any case, users may have privacy demands for all content types, so that a fine-grained differentiation may thus not be required in the context of this thesis.

Other research puts a special focus on emotional bindings between public displays and their users. Such emotions can be triggered by various stimuli, e.g., photos [76, 97, 128, 224], stories and texts [52, 131], music [132], drawings [57], or opinions [22, 215]. They all allow users

to shape their personality in the public. Creating emotional bindings may actually increase the attractiveness of public displays since a similar effect can be observed when looking at real (physical) assets:

Many people enjoy making a statement for others to see. They carefully choose clothes, jewelry, and footwear to indicate profession, interests, mood, or the music or subculture with which they identify. Houses, gardens, and cars are all accoutrements of social status and objects of conscious or unconscious display. Even graffiti can be seen as a particular form of personalization in which disaffected youths tag locations with their personal messages. Display networks might help channel this creativity and desire for personalization. —*Davies et al. [65]*

Though personalization may help to mitigate the effects of display blindness, it may not be sufficient to focus on this approach only. Müller et al., for example, discovered that the display location seems to have an even more significant influence on viewing times than the content itself [152]. Along the same lines, Davies et al. concede that though they presented a suitable technical solution to personalization, they failed to establish compelling use cases:

We do however note that we have not, to date, been able to develop or demonstrate compelling uses for the technology within the content of our public-display system. [...] [One] application was viewed most favorably by potential users but even for this application it is not clear that a large number of users would actually invest the time to make use of the system. —*Davies et al. [63]*

In an opposite approach, Greenberg et al. [86] analyze how *dark patterns* in the field of *proxemic* interaction may attract the users' attention to public displays. As the name implies, however, these dark patterns (sometimes also called *anti patterns*) should be avoided, since some users may not regard them as beneficial despite their apparent effectiveness. The pattern called "bait and switch" is an example:

The system baits the viewer with something that is (from the viewer's perspective) desirable, but the system then switches it to something else after the person directs his or her attention to it and moves closer. —Greenberg et al. [86]

In conclusion, public displays should comprise multiple approaches to mitigate the effects of display blindness in a multi-layered way. It should, however, not trick the user or involve deception. One particularly promising cornerstone appears to be personalization. Showing personalized content on a public display may, however, also raise some privacy concerns: Depending on the type of personalization, private or sensitive information could be shown and thus become public. The next section discusses this aspect in more detail.

3.2. Preserving Privacy During Personalization

Ostensibly, personalization and privacy on public displays may seem to be at odds with each other. This tension is also known as the *personalization-privacy paradox* [219]. Identifying specific reasons for this paradox, however, may be challenging. Section 6.2 introduces the notion of privacy in more depth. This section, however, underpins the general motivation behind this thesis. Looking at the definition of personalization may help to analyze the conflict:

per-son-al-ize. To design or tailor to meet an individual's specifications, needs, or preferences: "a personalized search engine;" "personalized learning." —*Dictionary.com Unabridged*¹

Depending on the individual user, a particular need or preference may be a very personal or even intimate piece of information. For example, large font sizes on public displays could indicate visual deficiencies or needs; automatically showing stock prices while approaching a display could reveal certain interests or preferences. The personalization-privacy paradox thus delineates a clash of interests: Unveiling personal information to systems may help to personalize those systems and make them more comfortable; at the same time, this comfort comes at the cost of decreased privacy.

This is, of course, just one facet of privacy as explained in Section 6.2. There are also other, more subtle aspects that need consideration in order to build privacy-preserving personalized public display systems. Davies et al. emphasize this in their research recommendations:

[...] The requirement to decide when personalized content can be presented without impacting on a viewer's privacy coupled with the need to manage the collection, storage and exchange of data to facilitate this represent fundamental challenges to the adoption of pervasive display networks.
—*Davies et al. [62]*

Though it seems reasonable to make the "collection, storage and exchange of data" [62] a transparent process to the users, this may not be possible in all situations. Some of these aspects may be very technical and laypeople may thus have a hard time understanding them. Users may have to trust systems that process their personal data to

¹personalization. Dictionary.com. Dictionary.com Unabridged. Random House, Inc. <http://dictionary.reference.com/browse/personalization>, accessed: March 04, 2015.

some degree. These systems should thus strive to preserve the user's trust, since "a wrong decision can negatively influence the user's acceptance of a system, cause frustration and, as a result, make users abandon the system" [241]. Similarly, Huang et al. [99] also expect responsible privacy policies to have a significant impact on the success of public display applications.

In 2013, the personalization-privacy paradox regained public interest fueled by the revelations of large-scale governmental surveillance. In the following months, extensive media coverage and public discussions raised the general privacy awareness. Yet, the public's interest started to decay over time, as it is the case for most political affairs. Nevertheless, the incident may have had a subconscious impact on the users' general attitude towards privacy. For example, users may have become more cautious about who actually possesses, i.e., stores, their data and information. Davies et al. already anticipated this issue one year earlier and proposed this design recommendation:

To maintain any semblance of privacy, an approach that requires users to register with some sort of central server that will subsequently track their movements in front of world-wide displays is not an option. Instead, users should stay in control of their data and decide for themselves when to personalize a particular display and what information to provide to support this. —*Davies et al. [65]*

4

Objectives

The previous chapters presented observable facts and phenomena, i.e., the proliferation of public displays, the effect of display blindness, and its psychological purpose of an information filter. Yet, there may also be negative aspects to display blindness, so that all stakeholders of public displays—operators as well as users—could benefit from mitigating these unfavorable characteristics.

Personalization appears to be a promising approach to address display blindness [65, 97, 108, 143]. Yet, pursuing the idea of personalizing public displays alone, especially without considering the users' demands for privacy, is not advisable [62, 65, 99, 241]. Thus, intrigued by the personalization-privacy paradox [219], this thesis explores a number of research questions as presented in Section 4.1. The answers to those research questions lead to the scientific contributions as summarized in Section 4.2. As already emphasized in the previous sections, the topics of privacy and personalization have a certain complexity to themselves, see also Section 6.2 and 6.3. Scrutinizing the interaction of both with respect to personalized public display systems may be even more complex. Section 4.3 thus outlines the scope of this thesis and demarcates aspects that have been excluded. Finally, Section 4.4 describes and visualizes the structure of this thesis.

4.1. Research Questions

To break down the complex matter of privacy-preserving personalized public display systems, this thesis subdivides the topic into three more manageable research questions RQ1–RQ3. Each of these research questions can be decomposed into a number of sub-questions, that are supposed to clarify the overall intention.

RQ1: What are main privacy threats on public displays?

- Is there a privacy threat model for public displays?
- To what extent are existing models applicable?
- Which application scenario requires the most privacy?

RQ2: What are countermeasures to those privacy threats?

- How can countermeasures be compared and differentiated?
- How to incorporate countermeasures into existing public displays?
- Do countermeasures impact the general public display usage?

RQ3: How to support the design process of public displays?

- What are common challenges to focus on?
- Can all steps be integrated in one process?
- Are Immersive Video Environments a valid methodology?

4.2. Scientific Contributions

This thesis provides results to the research questions in form of three tangible scientific contributions. They are presented in the remainder of this section. Besides these three contributions, this thesis comprises an extensive survey of recent literature on personalized public display systems, see Section 7.5. This survey uses a novel classification scheme to review 120 research projects on public displays. The scheme employs two established categorization dimensions: *user values* (application scenarios) and personalization usage models. The contribution of the survey provides insights in two domains: (i) the distribution of application scenarios and types of personalization being applied; and (ii) the distribution of privacy threats and proposed countermeasures. In both domains, the survey highlights apparent imbalances and points out areas that have been well covered by research as well as those that hold research opportunities.

CI: Privacy Threat Model

This contribution identifies key privacy threats that public displays may be subject to, and proposes a suitable threat model. A user study is conducted to explore what possible privacy threats there are on public displays. The results are mapped onto the *STRIDE* model [93]. CI then proposes a generic privacy threat model for interactive public displays based on an extended version of this model. This is an approach to a formal description of such systems, that can also be used to compare and analyze various characteristics. CI also contains a tangible prototypical tool that incorporates the findings, see Section 11.1. Researchers and designers may use this tool to design and evaluate

privacy-aware interactive systems. Furthermore, the tool is also used to evaluate C1, see Section 12.1.

C2: Countermeasures

Based on an extensive survey of 120 publications, C2 provides insights on countermeasures that may be applied to address some of the privacy threats identified by C1. Based on these findings, C2 presents a heat map of countermeasures, that visualizes the correlation between each countermeasure and all privacy threats. Researchers interested in privacy-preserving personalized public display systems as well as designers of such systems can use this heat map to quickly gain insights and draw conclusions. For example, researchers may focus on the “white spots,” as they point to combinations of privacy threats and countermeasures that have not been researched yet. Designers may use the heat map as a reference when designing, prototyping, or evaluating privacy-preserving personalized public display systems. Once they identified that their system may be subject to a particular privacy threat, they can look for “the most common countermeasure.”

To contribute to this set of countermeasures, C2 also contains three novel privacy-preserving approaches. *Visual multiplexing* allows for transferring multiple pieces of information, e.g., images, from public displays to smartphones, solely based on optical communication. This approach avoids conventional data networks, such as *WiFi* or *3G*, and may thus preserve the users’ privacy since no unique identifiers, e.g., MAC addresses, are transferred over a traceable connection. Subsection 10.2.1 presents more details on this countermeasure. *Visual highlighting* may be useful in situations, in which personalization is based on filtering screen contents. A flight departure board at an airport is an example: It might be sufficient to personalize this

large public display by highlighting one particular piece of information, i.e., the user's flight; the highlight is only shown on the user's smartphone and sensitive information, e.g., the final destination, may thus remain private. As with visual multiplexing, the data required to realize the visual highlights is also transferred optically from the public display to the smartphone. Subsection 10.2.2 discusses this approach in more depth. *Visual interaction* allows for communication in the opposite direction, i.e., from the user's smartphone to the public display. The approach supports adaptive user interfaces that can be tailored to suit a particular application scenario. Similar to the first two countermeasures, the necessary data is transferred based on optical communication only. Tracking users may thus be more difficult than in conventional (IP based) networks. Subsection 10.2.3 elaborates this approach further.

In order to provide concrete scientific contributions and allow for well-founded evaluations, see Section 12.2, each countermeasure has been realized as a prototypical implementation, see Section 11.2.

C3: Process Integration

Based on the desire to provide tangible scientific contributions, C3 comprises a methodology and tools to support the design, prototyping, and evaluation of privacy-preserving personalized public display systems, see Section 11.3. C3 proposes a novel method to engineer such systems based on realistic audiovisual simulations and a state-transition graph. C3 includes a systematic analysis of approaches to engineer public displays, a novel approach that integrates many of the benefits of previous approaches, an architecture for a toolkit implementing the approach, and an initial assessment of the approach

based on an example scenario and experiences from using it, see Section 12.3. Key benefits of this approach include high re-usability of simulated environments, reduced effort to simulate deployment sites and scenarios, as well as support for a broad range of prototypes, e.g., of varying fidelity, and design and evaluation methods. C3 can thus contribute towards simplifying and accelerating the development of privacy-preserving personalized public display systems.

4.3. Scope

This thesis explores the design of privacy-preserving personalized public display systems. The objective is to find answers to the aforementioned research questions and to transfer the outcomes into the scientific contributions presented in the previous section. To remain focused and concise, the breadth of this thesis had to be limited. Consequently, other intriguing research questions are beyond its scope. For example, the following aspects could be covered by future work in this research domain, as it is not the focus of this thesis to (i) invent new hardware, (ii) invent new evaluation methods, (iii) propose software tools ready for production, (iv) propose novel public display application scenarios, (v) propose a universal privacy threat countermeasure, (vi) propose a final definition of a threat model for public displays, (vii) propose a final definition of privacy, or to (viii) propose a final definition of context.

4.4. Thesis Outline

This thesis is structured in four parts as visualized in Figure 4.1. Part I introduces the overall topic, presents the evolution of public displays, and motivates the research objectives. Part II lays out the scientific foundation of this thesis. It presents the employed research methodology, introduces key concepts, e.g., definitions of context and privacy, and points to related work.

Part III constitutes the main body of this thesis: It addresses the overall topic of designing privacy-preserving personalized public display systems. This part covers the research questions R1–R3 and provides the scientific contributions C1–C3. Each scientific contribution is instantiated in a prototypical implementation, which is evaluated in turn, e.g., by expert reviews (C1, C3) or user studies (C2). Finally, Part IV reflects on all introduced approaches and contributions of Part III and discusses the results with regard to the research objectives declared in Part I as well as the scientific foundations of Part II. It also concludes the thesis by summarizing the main results as well as scientific contributions, and presents possible directions for future work.

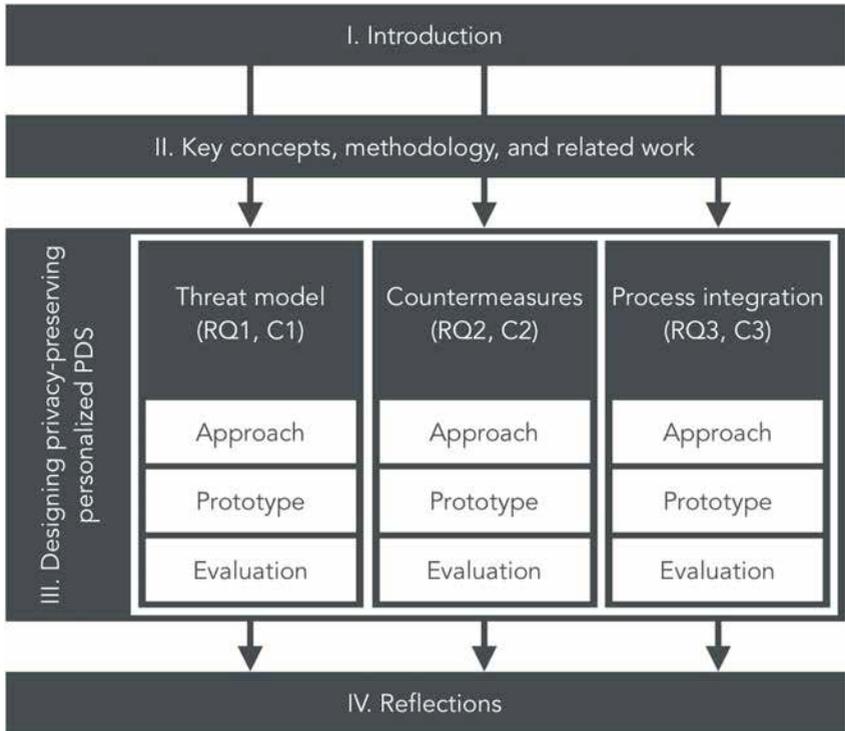


Figure 4.1.: Visual representation of the thesis structure. Part I introduces the topic and motivates the work. Part II provides the scientific foundation. Part III contains the research questions RQ1–RQ3 and the scientific contributions C1–C3 along with the corresponding approaches, prototypes, and evaluations. Finally, Part IV concludes with reflections of the results and an outlook of future work.

II

Methodology, Key Concepts, and Related Work

5

Research Methodology

In general, there are two widespread approaches to research, i.e., qualitative and quantitative methodologies. According to Alt et al. [12], research related to public displays mainly uses one of at least three types of studies for both methodologies: descriptive, experimental, and relational studies. However, as Alt et al. note, the relational approach is less frequently used in this domain, “because not many relationships between different dependent variables are considered to be interesting” [12]. Thus, this thesis focuses on the first two approaches and applies them as described in Table 5.1. The remainder of this chapter explains how descriptive studies were used in the course of this thesis, and how experimental studies were applied.

The first research question (RQ1) is concerned with the main privacy threats on public displays. Sections 2.2 and 3.1 point out, why there should be a special focus on the actual users of public displays: Systems should be tailored to suit the users’ needs in order to re-establish the users’ acceptance towards public displays. To address this special focus, the answers to RQ1 are based on a user survey. Once the threat model (C1) has been defined based on the results of the survey, the model is evaluated by expert interviews. Such interviews relate to the paradigm of “asking users,” as identified by Alt et al. [12]. The

Table 5.1.: Study types applied in this thesis as identified and explained by Alt et al. [12].

	Explanation	Application
Descriptive	Objective descriptions based on observations, interviews, focus groups, or logs	Threat model (C1): expert interview Countermeasures (C2): literature review Process integration (C3): expert interview
Experimental	Exploration of causality based on field or lab studies	Threat model (C1): user study, survey Countermeasures (C2): lab and user studies Process integration (C3): user study

outcomes of the expert interviews may thus indicate whether the expectations of the actual users were met. Observing the actual users during their interaction with the prototypical implementation would have been another way to verify this. Such observations, however, are beyond the scope of this thesis and thus constitute possible future work, see Section 15.2. The user survey also contained an experimental study, that analyzes whether the users' privacy demands correlate to different application scenarios. Evaluating such a causality is part of experimental research, as defined by Alt et al.[12].

The second research question (RQ2) strives to find countermeasures that may be applied to personalized public display systems in order to mitigate possible privacy threats. The list of countermeasures (C2) is the result of an extensive literature review that comprises 120 papers

in the domain of personalized public display systems. C2 is thus based on a descriptive observation. Other means of evaluation, e.g., expert interviews, would have required significantly more time, while possibly providing a less broad overview. Future work could, however, evaluate particular countermeasures based on individual expert interviews. The three novel countermeasures presented in Section 10.2 were evaluated in lab and user studies in order to assess and evaluate their usability and technical feasibility.

Finally, the third research question (RQ3) analyzes how the scientific contributions C1 and C2 could be incorporated and made available as concrete results. The outcome is an approach towards a process integration (C3). This approach comprises a new methodology, that may be used by experts, such as system designers or researchers, to design, prototype, and evaluate privacy-preserving personalized public display systems. Consequently, C3 was evaluated based on expert interviews and a user study.

6

Key Concepts in Public Display Systems

This chapter introduces key concepts, such as definitions of *context* and privacy, and defines their perception and usage in the context of this thesis. Some of the concepts are discussed controversially in literature. However, it is not the aim of this chapter to contribute yet another conception to this discussion, but rather to present an overview of common concepts that may help to work on the research questions and scientific contributions presented in Chapter 4. First, Section 6.1 presents different notions of context, an important term that is referred to in the subsequent discussion of privacy in Section 6.2. Next, Section 6.3 presents different approaches to personalization of public displays. Afterwards, Section 6.4 highlights common design concepts, e.g., design spaces, with regard to public display systems. Finally, Section 6.5 defines the terms threat, threat model, and countermeasure.

6.1. Context

It is important to define the meaning of context in the frame of reference in this thesis, as it seems to have a major impact on the particular privacy requirements of individual users: “Privacy [...] will be heavily dependent upon the context in which it occurs” as noted by O’Neill et al. [168]. For example, automatically showing photos of the user’s last vacation may be innocuous in the case that possible bystanders are friends or there are no bystanders at all. However, showing the same photos may be embarrassing or inappropriate if colleagues or business partners may see them, too.

The previous example indicates, that people and their corresponding roles may be one important dimension of a particular context. The definition of context presented by Dey may be more generic as it comprises other dimension as well:

Context is any information that can be used to characterise the situation of an entity. An entity is a person, place, or object that is considered relevant to the interaction between a user and an application, including the user and applications themselves. —*Dey [72]*

Dey’s definition seems reasonable and applicable to many theoretical deliberations. At the same time, however, it is rather general and less specific. Consequently, it may be too extensive in some situations. In terms of public displays, it remains unclear what “any information” would be that characterizes the system in particular. Thus Nissenbaum’s approach towards a definition of context with respect to privacy could be regarded as more concrete:

Contexts are structured social settings characterized by canonical activities, roles, relationships, power structures, norms (or rules), and internal values (goals, ends, purposes). [...] Contexts are not formally defined constructs, but [...] are intended as abstract representations of social structures experienced in daily life. —*Nissenbaum [161, pp. 132]*

Nissenbaum continues that her definition of context should not be regarded as final. She rather claims to “create a generalized snapshot of a context based on attributes observed across concrete instances, ultimately testable in the real world” [161, p. 134]. She also emphasizes that the perception of a particular context may be dependent on the country or geographic region, and may even vary between individual members of a society.

An important aspect of Nissenbaum’s perception of context is *nesting*: One context can be embedded or incorporated into another context. For example, consider the first scenario (“Getting to Work”) presented in Chapter 1: The public display in the entrance hall of the train station, that shows promotions, may belong to a context entitled “advertisement in train stations,” which may be a subcategory of “advertisement in public transport,” which may, in turn, belong to “general advertisement.” Nesting may thus help to generalize or to specify contexts and their related privacy requirements as necessary.

In her book, Nissenbaum establishes the notion of *context-relative informational norms* to underpin her framework of *contextual integrity* and her concept of privacy, see Section 6.2. Besides contexts, she identifies three parameters that appear to be key with respect to the structure of these norms: *actors*, *attributes*, and *transmission principles*.

There are three types of actors, i.e., “senders of information, recipients of information, and information subjects” [161, p. 141]. In the second example presented in Chapter 1 (“Getting Food”), the sender of the information would be the company that operates the canteen, the recipients would be all students or employees, and the information subject would be the available foods. Here, the subject is an inanimate object, i.e., a dish, but it could also be a human being in another scenario, e.g., a public display showing a patient’s vital parameters.

Nissenbaum’s idea of attributes relates to the type or nature of a particular information: “what it was about” [161, p. 143]. An attribute would thus—from a technical point of view—correlate to the bare data, e.g., the name or price of a dish. Yet, Nissenbaum refuses to present a more precise definition of attributes and rather settles for an intuitive sense.

Finally, the notion of transmission principles appears to be striking for Nissenbaum’s framework of contextual integrity. She defines the term as “a constraint on the flow (distribution, dissemination, and transmission) of information from party to party in a context” [161, p. 145]. Considering the shopping example presented in Chapter 1, it is common practice for supermarkets to analyze their customers’ shopping preferences via loyalty cards, e.g., “Payback” or “DeutschlandCard.” Though customers may not really appreciate this aggregation of data, they accept it and expect the operators to use the data for internal purposes only and to keep the data confidential. This expectation would be breached, if the public display in the example would use this aggregated data to advertise goods that particularly match an individual user. In this case, the information that is expected to be kept confidential is made public, by showing it on the display.

Overall, these dimensions to context as proposed by Nissenbaum may well suit the requirements for describing privacy-preserving public

display systems. A review of related work, see Chapter 7 as well as Sections 10.2 and 11.2, also suggests to add these aspects: location, form factors, *environmental factors*, interaction, *multi-display networks*, and *legal constraints*, as described in more detail in Chapters 9 and 10. The term location is also used to refer to the *spatial context* of public displays, that has been attributed to be of particular importance [73, 178].

As mentioned above, there may be more approaches towards a definition of context in the literature, for example, the one proposed by Zimmermann et al. [247]. However, the aim of this section was not to present a final rationale, but to open the mind of the reader for this complex aspect with regard to designing privacy-preserving personalized public display systems.

6.2. Privacy

To define privacy in a canonical way is challenging. Nissenbaum postulates that “one point on which there seems to be near-unanimous agreement is that privacy is a messy and complex subject” [161, p. 67]. Along the same lines, Langheinrich claims that “privacy is related to, but not identical with: secrecy, solitude, liberty, autonomy, freedom, intimacy, and personhood” [125]. Another closely related term to privacy is *security*. While both, privacy and security, are sometimes confused or used as synonyms, drawing a sharp line between them appears to be a difficult task as well since “security is an integral part of any privacy solution” [125]. Despite this apparent complexity, most scholars agree, that privacy cannot be described adequately by a one-dimensional spectrum ranging from public to private. Nissenbaum refers to that notion as the *public/private dichotomy* [161, p. 89], which she tries to disassemble in the course of her book [161, p. 232]. O’Neill et al. also postulate that “any essentialist public/private dichotomy is

over-simplistic” [168]. This thesis thus draws from more nuanced approaches to privacy, as presented in the remainder of this section.

The appearance of technical devices sometimes initiated discussions about possible (re-)definitions of privacy, as indicated by Langheinrich [125]. For example, around 1890, the emergence of the hand-held photo camera sparked an approach towards such a definition. Apparently, people were concerned about the impact of this device on their everyday lives, as it was now possible for anyone to capture certain moments, e.g., people walking down the street or attending gatherings, and preserve these moments in an objective way—regardless of whether the people on the photograph would approve. One prominent definition of privacy at that time was conceived by Warren and Brandeis that defines privacy as “the right to be let alone” [236].

Westin interprets privacy as “the claim of individuals, groups, or institutions to determine for themselves when, how, and to what extent information about them is communicated to others” [238]. His definition may be motivated by the advent of large mainframe computers in the 1960s, as these computers were able to automatically evaluate large amounts of data in new, previously unknown ways.

Such a change in expected or common behavior—often related to technological advances—is a central aspect of Nissenbaum’s approach to privacy introduced in the framework of contextual integrity. Nissenbaum’s notion of transmission principles, see Section 6.1, confines the flow of information in a particular situation. If, for whatever reason, information appears to flow unexpectedly, i.e., against established transmission principles or expectations, the involved parties may regard this as a violation of their privacy. Langheinrich seconds this:

Many people wish to control the flow of information about themselves [...]. Privacy is also about opportunistic data use (i.e., data that has been given for one purpose is “recycled” for another) or involuntary disclosures (e.g., someone who is entitled to receive information in general should not have gotten this information in a particular, unexpected situation). —*Langheinrich [125]*

In an attempt to structure possible “harms and problems” [212] of privacy breaches, Solove proposed a taxonomy of privacy “to guide the law toward a more coherent understanding of privacy and to serve as a framework for the future development of the field of privacy law” [212]. This indicates that Solove’s point of view is that of a jurist rather than the one of a researcher or system designer. Nevertheless, Solove’s taxonomy may be useful to these people as well, as it names concrete privacy issues to focus on. Table 6.1 shows Solove’s four general categories along with their sixteen subcategories.

The first three of the four general categories seem to have an apparent relation to privacy on public displays, for example: Users might feel under constant surveillance if pervasive public displays would be able to track them, e.g., via cameras. It may also be desirable to make the data aggregation and identification of users a transparent process. The most challenging issue with regard to privacy on public displays might be cases of breached confidentiality or information disclosure in general, as the (large) screens may increase the accessibility of personal information. The suitability of the last general category may be less apparent in the context of this thesis since invasions “do not always involve information” [212], but are “direct interferences with the individual, such as intruding into her life or regulating the kinds of decisions she can make about her life” [213]. However, Solove also claims that “spam, junk mail, [...] and telemarketing are disruptive in a sim-

Table 6.1.: Solove's taxonomy of privacy [212].

General category	Subcategory
Information collection	Surveillance Interrogation
Information processing	Aggregation Identification Insecurity Secondary Use Exclusion
Information dissemination	Breach of confidentiality Disclosure Exposure Increased accessibility Blackmail Appropriation Distortion
Invasion	Intrusion Decisional interference

ilar way, as they sap people's time and attention and interrupt their activities" [212]. This, however, may relate to public displays, especially in terms of pervasive advertising and the dystopian prospect of commercials depicted in the movie "Minority Report," see Section 2.1.

Both, Solove and Nissenbaum agree that privacy may not be defined as a sheer technical term, but as a notion deeply rooted in the society it is used in: "Privacy cannot be understood independently from society" [212], or as Nissenbaum puts it:

According to the framework [of contextual integrity], finely calibrated systems of social norms, or rules, govern the flow of personal information in distinct social contexts [...]. These norms [...] define and sustain essential activities and key relationships and interests, [and] protect people and groups against harm [...]. —*Nissenbaum [161, p. 3]*

Apparently, there is a large interest in privacy “spanning disciplines from philosophy to political science, political and legal theory, media and information studies, and, increasingly, computer science and engineering” [161, p. 67]. In stark contrast to this, however, the general public does not seem to ascribe importance to privacy in equal measure. Section 7.2 presents evidence for this behavior based on scientific results. Solove addresses the public’s attitude in his article about the common “I’ve Got Nothing to Hide” argument [213]. According to Solove, this belief is rooted in the aberrant understanding of people that privacy is all about hiding something, e.g., bad habits, sexual preferences, crimes, or other unlawful activities. But privacy is more than that: It is the “plurality of privacy problems implicated by government data collection and use beyond surveillance and disclosure” [213]. While referring to Jeffrey Reiman, Nissenbaum lists four types of risks that may occur if people would be truly deprived of privacy: risks of extrinsic and intrinsic losses of freedom, symbolic risks, and risks of ‘psycho-political metamorphosis’ [161, p. 75]:

Extrinsic losses of freedom occur when people curtail outward behaviors that might be unpopular, unusual, or unconventional because they fear tangible or intangible reprisals, such as ridicule, loss of a job, or denial of benefits. Intrinsic losses of freedom are the result of internal censorship caused by awareness that one's every action is being noted and recorded. [...] Those being watched [...] are thus deprived of spontaneity and full agency as they self-consciously formulate plans and actions from this third party perspective. [...] The symbolic risk of institutional structures that deny individuals the capacity to withdraw is that they deny them this expression of self-ownership.³ The fourth risk [...] [is] that if people are subjected to constant surveillance, they will be stunted not only in how they act, but in how they think. They will aspire to a middle-of-the-road conventionality—to seek in their thoughts a “happy medium.” —*Nissenbaum [161, pp. 75-76]*

In summary, privacy is substantial component of people's everyday life, because it allows them to act and evolve freely. People's behavior would change if they lost all their privacy, even though some would deny such change and willingly relinquish their privacy in order to pursue a higher goal, e.g., the fight against terror [213]. With the ongoing proliferation of public displays throughout urban environments, privacy needs to be re-considered, as transmission principles, i.e., generally accepted flows of information, may be about to change or they already have. Thus, the research community as well as designers and users of personalized public display systems may benefit from an overview of existing application scenarios (Section 7.3), possible privacy threats (Section 10.1), applicable countermeasures (Section 10.2), and uncovered areas remaining for future work (Section 7.5).

6.3. Personalization

As discussed in Section 3.2, the verb “to personalize” means to adjust something to the specifications, needs, or preferences of an individual. Since this is a rather general definition with regard to public displays, the term should be specified further. First of all, some publications refer to personalization as *user modeling* [5, 149, 155, 216, 228]. Furthermore, the point of view, see Section 2.2, should be considered: Display operators may without doubt personalize the outer appearance of public displays according to their needs, e.g., to match their corporate identity. They may also decide on where public displays should be installed at, i.e., they may define their location. When talking about public displays, however, the term personalization does most often not refer to the operator’s point of view, but to the user’s point of view. From the user’s perspective, certain attributes of a public display are usually immutable, e.g., the location, or the shape and size—i.e., the physical attributes. In contrast, virtual attributes, foremost the display content, may sometimes be changed or manipulated by the users. The scope of personalization ranges from general pre-selected content to individual personal information. Section 7.3 presents some projects that pursue the approach of personalizing public displays. Various sources can be exploited to obtain personalized information or personalized data. For example, one strand of research analyzes the use of social networks [9, 96, 107], while another strand focuses on other means, e.g., smartphones [97, 190, 203]. Another aspect worth consideration is the type of personalization. Davies et al. [64] introduced three personalization usage models: walk-by, longitudinal, and active personalization. Section 7.3 presents this notion in more detail. In this thesis, personalized public displays are systems, that can be adjusted to the specifications, needs, or preferences of the users—in terms of viewers.

6.4. Design

Public display systems consist of at least two components: hardware and software. The hardware is defined, for example, by the actual screen, protective casing, or various interfaces, such as keyboards or cameras. The software is the second vital component, as it includes, e.g., the content scheduling algorithm or the privacy-preserving logic that protects the user’s sensitive information. Nowadays, of course, the majority of public displays also require a network connection and some type of *backend* that provides the actual contents. This part of the system architecture, however, is very similar to other existing services, e.g., customized RSS news feeds. Thus, this thesis mentions that part only on the sidelines.

When looking at the hardware and software components of public displays, there are established design processes for each of them. The discipline of software engineering, for instance, concentrates on “the application of a systematic, disciplined, quantifiable approach to the development, operation, and maintenance of software” [1, p. 67]. To achieve this goal, various models have been defined and extended over the years. One of them is the *spiral model* by Boehm [29] introduced in 1986. Figure 6.1 shows the model. There are four quadrants: (i) determine objectives, (ii) identify and resolve risks, (iii) development and test, and (iv) plan the next iterations. This thesis puts a special focus on the second and third one, as its scientific contributions help to identify (C1) and resolve (C2) privacy threats and present ways to handle the development and test via integrated processes (C3).

A common saying is, however, that “the whole is greater than the sum of its parts,” which is most often credited to Aristotle. With regard to public displays, this saying could be interpreted in a way that designing sturdy hardware and robust software separately from each

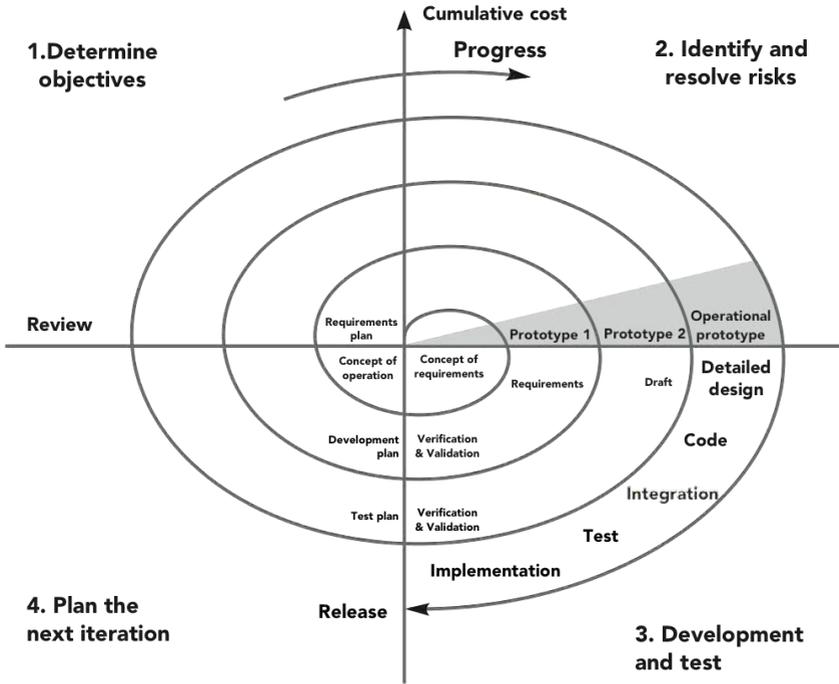


Figure 6.1.: The spiral model by Boehm [29]. Graphic by [59].

other would not be sufficient in order to build a successful privacy-preserving personalized public display system. As explained in Part III, many aspects contribute to the design of such systems, many overarching the hardware-software boundary.

One of the first public displays, called “Hole-In-Space,” was setup in 1980 [12]. Alt et al. [12] state that despite a long-lived interest of science in this topic, no commonly accepted design guidelines emerged. Rather, there exists an amalgamation of approaches, that have been applied in individual research projects. Thus, Alt et al. compiled a list of these approaches along with common research questions, that the approaches address. Table 6.2 summarizes their results by indicating combinations of research paradigms and research questions commonly found in literature.

Table 6.2.: Combinations of research paradigms and research questions according to Alt et al. [12]. The ‘x’ indicates that this combination has been addressed in at least one publication. DBR stands for deployment-based research.

Research question	Research paradigm				
	Asking users	Ethnography	Lab study	Field study	DBR
Audience behavior	x	x	x	x	x
User experience			x	x	x
User acceptance	x	x	x	x	
User performance			x	x	x
Effectiveness	x	x	x	x	
Privacy			x	x	
Social impact		x		x	

Apparently, field and lab studies seem to be well suited to provide answers to most research questions. Regardless of the ongoing discussion about varying validities or reliabilities of field-based vs. lab-based results, this thesis strives to ease the gap between both approaches: The scientific contribution C3 proposes a way to merge the advantages of both approaches in order to create reproducible, high-fidelity simulations of interactive environments, see Sections 10.3 and 12.3.

In any case, each paradigm requires a concrete research method in order to provide actual results. In this regard, Alt et al. [12] compiled the following list: (i) interviews, (ii) questionnaires, (iii) focus groups, (iv) observations, and (v) logging. Furthermore, (vi) rapid prototyping, (vii) literature reviews, as well as (viii) models and replicas can be added to this list [174]. In particular, the scientific contribution C3 supports the methods of logging, rapid prototyping, and models and replicas, see Sections 10.3 and 11.3.

6.5. Threats, Threat Models, and Countermeasures

To understand each of the terms introduced in this section, it is advisable to look at them in consecutive order as implied by the title of the section. The word threat is probably well-known and unambiguous in non-technical domains, as defined as follows, for example:

threat. A person or thing that is regarded as dangerous or likely to inflict pain or misery. —*Dictionary.com Unabridged*¹

¹threat. Dictionary.com. Collins English Dictionary — Complete & Unabridged 10th Edition. HarperCollins Publishers. <http://dictionary.reference.com/browse/threat>, accessed: March 10, 2015.

However, there are also technical definitions and standards that try to be more specific or more general, depending on their origin and purpose. The German Federal Office for Information Security (BSI), for example, interprets the term as follows:

[...] [A] threat is a condition or an event which can negatively affect the availability, integrity or the confidentiality of information, which in turn results in damage to the owner of the information. Basic threats can result from the effects of force majeure, organisational shortcomings, human errors, technical failure or deliberate acts.

—*German Federal Office for Information Security [84, p. 39]*

Overall, a threat is understood as something that is not desirable, and that may stem from a plethora of origins. The latter characteristic in particular renders comprehensive assessments challenging, for example: Who may threaten whom? Since technical threats play an important role in today's society in terms of IT security, engineers are striving for systematic approaches towards assessing threats. One particular approach to this are threat models. Threat models are a standard tool in software engineering. They are used to describe threats and attackers to a software system:

[A threat model] refers to a systematic review of a system design to discover and correct security problems at design-level. It is used to understand a product's threat environment and defend against potential attacks. Threat modeling allows methodically identifying, evaluating and rating application threats and vulnerabilities. By rating threats, one can address threats with suitable countermeasures in an order, starting with the threats that have greatest risk.

—*Kaur and Kaur [111]*

Since software can take many different forms, it is very challenging to cover all possible threats. However, there are (categories of) threats that can be applied to most software systems. Hernan et al. [93] proposed the STRIDE threat model that identifies common threat categories and consolidates them in a single model. Each letter of the acronym corresponds to a category of threats that a system may be subject to, as explained below. Section 10.1 presents the rationale for choosing STRIDE as the base for a privacy threat model for public displays (C1). It also discusses alternative theoretical groundings and proposes the the final design of C1.

As part of the scientific contribution C1, this thesis identifies one further privacy-related aspect of relevance for public displays: *decontextualization*. This result is based on an extensive review of related work (120 papers about personalized public displays published between 2000 and 2014, see Section 7.5). The remainder of this thesis thus refers to the STRIDE model, that has been extended by decontextualization, as the *STRIDED** model.

Spoofing. In a *spoofing* attack, attackers pretend to be someone else—they spoof their identity. Shoulder-surfing is one way to obtain required credentials, e.g., passwords, which can be performed quite easily in many usage scenarios for public displays, for example, at outdoor ATMs [68]. Consequently, spoofing is broadly discussed in the literature. Though the main focus is on spoofing the users' identity, this threat also exist in the reverse way: Compromised public display systems might pretend to be an authority the user can trust.

Tampering. *Tampering* is a rather general threat, which almost all systems with public interfaces are vulnerable for. An example for a tampered public display is a manipulated ATM, that is misused to spy out the PINs of customers. This attack could eas-

ily be extended to other public display applications in general, for example, when users have to input their password via forged keyboards. In addition to attacks that target the input on public displays, tampering can also be used to manipulate the output of information, for example, by obscuring parts of the screen or injecting non-curated content.

Repudiation. Kohnfelder and Garg define *repudiation* as “an untrusted user performing an illegal operation without the ability to be traced. [...] [Repudiation] threats are associated with users (malicious or otherwise) who can deny a wrongdoing without any way to prove otherwise” [116]. Users could thus deny using the public display for receiving or entering information, e.g., changes to an itinerary or illegal contents. Vice versa, the public displays could deny the reception or the presentation of information, e.g., receiving a file or showing wrong flight departure times.

Information Disclosure. If the users’ private or sensitive information becomes publicly available without their consent, this is called *information disclosure*. This threat can have various manifestations and is thus extensively covered by related work. According to Davies et al. [64], there are at least three different types of information disclosure: (i) *location disclosure*, (ii) *content disclosure*, and (iii) *use of display infrastructure*. The first type (i) addresses the threat of tracking the users’ position while using a mobile device with a public display, for example, via tracking the MAC addresses of the smartphones or (at the side of the mobile service provider) tracking the GSM cell towers. An example for the second type (ii) would be users browsing their calendar without noticing that another user is approaching. As a result, the other user could get a glimpse of, e.g., possibly sensitive calendar items. Finally, the third type (iii) is an amalgamation of the previous

types, yet in an even wider scope. Only by using a personalized public display system, e.g., explicit logins via username and password or implicit logins via proxemics, the identity and *presence* of users can be tracked. A well-known example of this threat is depicted in the dystopian view presented in the movie “Minority Report.” People can be identified by public displays (exemplified in the movie by advertising boards) even if they do not approve of this. The third type of information disclosure is also comparable to the *telescreens* described in George Orwell’s book “1984.”

Denial of Service A *denial of service* (DoS) attack usually floods or jams the target with an excessive amount of (network) data. However, the definition of this attack can also be broadened to include the contents of public displays: (i) Attackers could block the system by claiming interaction for a prolonged period of time so that others cannot use the system anymore. (ii) The service offered by the public display could be “flooded” with user generated content, so that real twitter feeds could be usurped by fake messages, for example. (iii) A denial of service attack can also threaten users’ privacy more directly, for example, if the system ceases to respond while users are interacting with it. The screen could then freeze while showing sensitive personal information, and the information would be publicly visible for an extended period of time without users being able to remove it.

Elevation of Privilege. This category can be mapped to the *security property* authorization, as suggested by Youn et al. [246]. *Elevation of privilege* is sometimes also called escalation of privilege. It includes all types of attacks that exploit faulty software or design flaws in general to gain access to information or services that the owner did not provide access to. Besides the abuse of software or hardware errors, this attack can be extended, e.g., to include

social engineering: Attackers could try to gain the victim's trust in order to obtain access to private or sensitive information.

Decontextualization. Information is always embedded in a certain context [183] and may become meaningless, if the corresponding context cannot be deduced. The lack of context leaves the interpretation up to the individual recipient, and may result in various diverging conclusions [37]. Often, the information shown on public displays can be perceived by the active user, but also by *passersby*, who may (consciously or unconsciously) draw conclusions about the perceived information and the particular user. These conclusions may be to the users' advantage, but also to their detriment. For example, a couple uses a public display in front of a travel agency to look at offers. Though they cannot afford the high priced trips, they look at the offers anyway. Passersby do not know that and may conclude that the couple is quite wealthy. In contrast, however, a male user checks his e-mails on another public display. He could stumble across some spam e-mails, that contain adult content. Passersby do not know that those pictures are unwanted spam and may conclude that the user actually requested the content.

Decontextualization can also occur in other application scenarios. Li et al. [129], for example, report on a user study about mobile projector-based interfaces for indoor navigation. The results of their study show that some users seem uncomfortable when receiving navigational instructions on a publicly projected screen. Apparently, some users fear that passersby could draw (wrong) conclusions, for example, by inferring information from the projected destination.

The term countermeasure refers to all measures that can be put in place (i) to avoid a threat being applied in an actual attack (referred to as an *applied threat* [84, p. 38]) or (ii) to counter the effects of an applied threat. Countermeasures also mitigate possible chances of successful applications and its effects. Sometimes, such means are also called *safeguards* [84] or *security controls* [227]. In most cases, it is yet impossible to comprehensively identify all potential threats a system may be subject to. Thus, it is even harder to compile an extensive list of countermeasures. Furthermore, threats and corresponding countermeasures are always linked to a very specific context, see Section 6.1. Hence, it is difficult to provide specific design recommendation to system designers and researchers with regard to threats and countermeasures. Such recommendations are limited to very generic considerations, for example, as pursued by the IT-Grundschutz Catalogues published by the German Federal Office for Information Security [84] or the Open Web Application Security Project (*OWASP*) [227].

Ultimately, what constitutes a threat, a threat model, and a countermeasure to a privacy-preserving personalized public display is highly dependent on the very context of that system. As discussed in Section 6.1, a plethora of attributes may contribute to the context, comprising, for example, users, locations, contents, means of interaction, form factors, environmental factors, or legal constraints. Designing a “secure” and privacy-preserving system thus clearly appears to be a sophisticated task. This thesis thus strives to contribute to this challenge in three ways: (i) by providing a threat model (C1) that can be used by system designers as well as researchers to systematically assess major privacy threats; (ii) by presenting a list of common countermeasures (C2) that may mitigate some privacy risks; (iii) by proposing means to integrate the findings presented in this thesis into established processes (C3).

7

Related Work on Public Display Systems

The design of privacy-preserving personalized public display systems draws on several research areas. Chapter 6 introduced five important key concepts, that also help to structure the related work correspondingly: First, Section 7.1 presents related work on *context-aware* public display systems. The next Section 7.2 introduces publications with regard to privacy on public displays. Then, Section 7.3 covers related work concerned with personalizing such systems. Section 7.4 presents approaches towards the design of public displays, e.g., by pointing to design spaces and classification schemes. The subsequent Section 7.5 is dedicated to related work addressing possible threats, threat models, and countermeasures on public displays. The same section also presents the results of an extensive literature review, that points to research opportunities and contributes to the list of countermeasures (C2). Finally, Section 7.6 highlights some existing toolkits and frameworks, that designers and researchers of public displays may use at various development stages. As privacy, threats, and countermeasures are of special interest in the context of this thesis, the corresponding Section 7.5 covers this topic extensively, while the remaining sections highlight important related work briefly and succinctly.

7.1. Context

According to Dey, “a system is context-aware if it uses context to provide relevant information and/or services to the user, where relevancy depends on the user’s task” [72]. Some terms in this definition might require further explanation—one of them is context, see Section 6.1. Numerous publications use such terms as context or context-aware. Unfortunately, though, some of them do not define their notion of context, cf. [136, 138]. Context may, for example, relate to the content and its viewers [196] or the public display itself [148]. Consequently, these particular publications can only be located vaguely within the field of context-aware public display systems. Yet, the intention of this section is to point out the breadth of approaches towards context-aware public displays by highlighting a few particular publications.

Cardoso and José [44] introduce a framework for context-aware adaptation. The authors understand a context-aware public display to be “able to deliver ‘the right information at the right time’” [44]. They claim that most current public displays are context-related in terms of installation site and expected audience only. Yet, basing context-awareness on these two static attributes may not cater for the needs of highly dynamic public places: “In order to be efficient, digital displays need to target their audience’s needs, expectations and tastes” [44]. Their approach to target these needs is based on mapping *digital footprints* to context-aware contents. Such footprints can be any traces that users create while they interact with public displays—either implicitly, e.g., by passing it or looking at it, or explicitly, e.g., by pressing buttons. Along the lines of Cardoso and José, digital footprints—and thus the contexts of public displays—comprise: presence, *presence self-exposure* (e.g., profiles maintained by the users), *user suggested content*, and *actionable* (e.g., downloading, controlling, or rating contents).

The public display presented by Kurdyukova et al. [122] pursues an approach that focuses on social aspects, i.e., the “group context.” In the authors’ believe, it is unrealistic that public displays could accommodate the needs of individuals with regard to content, for example. They thus suggest to rather analyze groups of viewers in front of the public displays and show content depending on, for example, the group size, gender distribution, or apparent social inter-relations, e.g., whether individual members appear to be friends, acquaintances, or strangers to each other. Furthermore, their system is able to detect the emotions of the viewers and tag the shown content accordingly. This way, the content is supposed to please particular social groups after a certain period of time. The approach presented by Kurdyukova et al. thus employs the longitudinal personalization usage model as introduced by Davies et al. [64].

Wißner et al. [241] extended the approach by Kurdyukova et al. [122]. Wißner et al. also try to infer the social context of public displays. Their definition of social context comprises the gender of the viewers, whether viewers are arriving or leaving, the availability of mobile devices, and two more complex attributes: the social inter-relations, defined as alone, friend, acquaintance, and stranger, as well as the privacy level of the shown content, i.e., private and not private. Wißner et al. point out that their approach may raise “issues with user trust” [122], as it is vital that users trust in such adaptive systems and feel comfortable while using them. If trust cannot be established and maintained, the user’s acceptance towards public displays may wane: “There is an enormous need for sophisticated trust management in ubiquitous display environments in order to ensure that such environments will find acceptance among users” [241].

The *PriCal* system introduced by Schaub et al. [196] applies “context-aware privacy” to prevent onlookers from glancing at users’ sensi-

tive calendars. PriCal is also based on an analysis of the social context, although this analysis considers less attributes than the previous approaches: “Present persons are the main dynamic context feature [...]” [196]. Although this notion of context appears quite limited (with regard to the inherent complexity as discussed in Section 6.1), the authors acknowledge that “discriminant context features need to be tailored to the respective application. [...] Robust context detection and adaptation is important for users to entrust such systems with their personal information” [196]. In this regard, Schaub et al. seem to be in complete agreement with Wißner et al. [241].

While the previously presented publications interpreted context in particular with respect to group structures or shown content, Vande Moere et al. focus on the “[...] economic and urban context [...]” [148] that public displays are embedded in. Their work is concerned about large-scale *media facades*, a particular type of public displays, and thus touches upon various aspects of urban planning. Their understanding of context is described as a threefold concept: “that what is *in front of*, *on* and *behind* [emphasis added] the public display [...]” [148]. Vande Moere et al. suggest to make media architecture more responsive to changes of context, for example, when new buildings are erected in the vicinity of the public display system, which might impact the visibility of the system.

7.2. Privacy

There exists a large body of research on public displays and privacy. This might not only indicate a particular interest of researchers, but also a certain societal significance, see Section 6.2. The remainder of this section may thus not provide a fully comprehensive overview

of related work in this domain, but it highlights some characteristic publications: Firstly, there is a presentation about concepts that aim at improving privacy in general. Secondly, there is an introduction to projects, that apply some of these concepts to specific application scenarios. Next, there is a report on the commonly observed fact that—despite the apparent importance of privacy—users do not seem to value it much. Finally, the section concludes that a privacy threat model (C1) may contribute to the research community and the general public.

Brudy et al. introduce three methods to raise the user’s “awareness of shoulder-surfing moments” [40] while using displays in semi-public settings. The system informs the user about passersby who might cast a glance towards the display by either showing a border around the entire content, or showing a little 3D model representing the position and body orientation of the passerby at the bottom of the screen, or by indicating the particular screen area that the passersby is probably looking at. Besides these three approaches to raise the user’s awareness, Brudy et al. also propose three protection measures, or safeguards as they call them. The first approach lets the user move all screen areas that contain sensitive information right in front of him, so that the information is shielded by the user’s body. To do that, it is sufficient to perform a quick and unobtrusive gesture. Alternatively, the first approach lets the user hide all sensitive information by turning away from the display. While the first approach requires the user to explicitly trigger an action, the second approach is more implicit. The system uses a heuristic to identify personal and public display contents. Whenever passerby are identified in the vicinity of the display, the private parts are automatically hidden. An advantage of this approach may be that the protection takes effect immediately and does not require the user to react. However, a disadvantage could be that the information the user is currently working on is automat-

ically hidden. This may interrupt the user's workflow. To minimize this interruption, Brudy et al. propose their concept of silhouette protection. This way, the system only hides the parts of the screen that are not shielded by the user's body, while the areas directly in front of the user remain visible.

The *SPIROS* system introduced by Röcker et al. [185] does not focus on semi-public settings, but on home environments. They introduce the concept of a "private space" that spans the area around the user and the public display. As soon as another person enters that private space, the system tries to identify that person and its social relationship to the user. The system then uses this information to adapt the visible content automatically. *SPIROS* is thus comparable to the second approach presented by Brudy et al. Additionally, however, *SPIROS* also considers the relationship between the persons in front of the display. Röcker et al. refer to this approach as "context-adapted privacy protection," which is based on five levels of privacy: (i) "[...] private content [...] [that] is meant for eyes of its owner only;" (ii) "[...] information which is meant for intimate circle of persons only;" (iii) "[...] information which is family-internal and potentially accessible by all family members;" (iv) "[...] family-internal information [is hidden] [...], but [...] access to other personal information [is still allowed] [...]" [185]. Similar to Brudy et al., *SPIROS* employs a heuristic to identify sensitive documents and applications, e.g., based on keywords. *SPIROS* also proposes a set of countermeasures akin to the approaches presented by Brudy et al.: The system may display a message that informs the user about possible privacy breaches or cover the sensitive information; additionally, based on the desktop metaphor of popular operating systems at that time, *SPIROS* can minimize, hide, or close all windows containing sensitive information. A year later, the same authors evaluated their system in a formative user study [186]. Their results indicate that "users are in general willing to trust system-

based protection mechanisms [...]. The proposed combination of pre-defined privacy profiles [author's note: presumably the countermeasures mentioned above] and context-adapted information visualization proofed [sic] to be a good trade-off between usability and adequate privacy protection" [186].

Tacita [64] envisions the use of pervasive display networks and mobile devices to provide privacy on personalized public displays. The system is especially characterized by the fact, that no user profiles are stored within the infrastructure of the system. In contrast, the user's preferences are administered on the user's mobile device. To do that, the mobile device monitors the user's geospatial position via GPS. The device uses this information to retrieve a list of all available public displays in the user's vicinity. In a subsequent step, the mobile device requests all displays to provide a list of personalizable aspects, e.g., cloud-based applications and contents. Based on these lists, the device computes the individual—personal—display assets to be shown. Due to this architecture, users do not have "to have any form of trust relationship with displays that they encounter in their travels" [64]. Still, users have to trust the underlying (network) infrastructure and the software on their mobile devices.

While the approaches presented above strive to propose generally applicable privacy concepts, the following references to related work exemplify concrete applications and use cases. Early work in this domain was executed by Shoemaker and Inkpen [205]. They introduced the concept of *single display privacyware*, which they named in analogy to the existing research on single display groupware. The difference between both strands of research is, that the latter one assumes that users share all information on a public display in order to work on a specific task. However, in application scenarios related to single display privacyware, as presented by Shoemaker and Inkpen, users may

prefer that certain pieces of information are only visible to them for at least three reasons: (i) to save limited screen real-estate; (ii) to reduce the amount of available information to a minimum (cf. the *awareness overload* problem); (iii) to retrieve additional context-related data. The system is based on shutter glasses, that are synchronized to the frame sequence of the actual display. Thus, users can only perceive frames that contain their individual data. Due to this hardware setup, however, the system may not be suited to actually protect a user's sensitive information from other persons, as they could simply take off their glasses and see all the information. This apparent caveat may be dissolved, however, by looking at the definition of privacy used by Shoemaker and Inkpen: "Privacy, limiting the availability of information to a single user, serves to reduce the level of group awareness" [205]. Hence, their focus is not on protecting highly private information, such as bank account information, but on increasing the efficiency of collaborative work on public displays.

Berger et al. [26] introduced *symbiotic displays*, that consist of public displays and accompanying personal devices. The displays can be used to show anonymized personal information, e.g., e-mails, while sensitive information, such as names and dates, are blurred. The display of the personal device unveils the anonymized content. Thus, as opposed to the approach by Shoemaker and Inkpen [205], symbiotic displays may actually be used to protect private information from other users. To let the system automatically detect and blur pieces of information, Berger et al. propose three levels of sensitivity: (i) maximum level of sensitivity, (ii) medium level of sensitivity, and (iii) full text. In their prototypical implementation, "the majority of the blurred words correspond to names, dates, locals, and numbers" [26]. They also point out that "the selection of the words being blurred was performed manually," [26] as "marking content with sensitivity levels may be difficult" [26]. They also suggest a number

of approaches to facilitate this, such as analyzing the text grammar structure or scanning for words that are not commonly used. Nevertheless, this essential part of the symbiotic display system remains a challenging task.

PriCal is designed to be “an ambient calendar display that shows a user’s schedule similar to a paper wall calendar” [196]. Comparable to SPIROS, PriCal aims at providing context-adaptive privacy based on sensing all persons in the proximity of the display. Using various sensors, the system detects all persons entering an office that is equipped with PriCal. Based on a “hide then reveal strategy,” the system removes all sensitive data from the public display as soon as new persons appear. After the system was able to identify the person that just appeared, it reveals appropriate sensitive data accordingly, i.e., based on the (social) relation between the user and the person. Users may individually specify relations and the corresponding information to be shown or hidden. Akin to Tacita, PriCal also avoids to store the user’s preferences on a centralized server in order to create a *trust relationship* between the users and the system. In contrast to Tacita, however, PriCal keeps the individual information within the systems infrastructure, i.e., on the individual public displays installed in the user’s office. The results of their user study imply “that context-adaptive privacy mechanisms are perceived as useable and useful [...]” [196].

Baldauf et al. [22] analyze the user’s privacy requirements for *interactive opinion polls* on public displays. It seems natural that there may be special needs for privacy when casting a ballot, especially if applied in official votes or elections. Baldauf et al. present results of a user study that analyzes two dimensions: the voting technique (“public touch interface, personal smartphone by scanning a QR code, from remote through a short Web address” [22]) and the type of poll question (“general, personal, local” [22]). Their results indicate, that people

prefer to vote publicly, i.e., by using a touchscreen, so that their vote can be seen by others, which may spark social discourse. Apparently, the type of question does not have an influence on the favored voting technique, but on the overall number of participants.

Similar to Baldauf et al., De Luca and Frauendienst [67] are also interested in input methods for devices in public environments. They claim that users are often required to enter information on public terminals, e.g., public displays. Using a keyboard, for example, bears the risk that passersby might acquire sensitive information via shoulder-surfing. They thus propose a system called *PocketPIN*, which lets users input sensitive information, such as passwords, via personal mobile devices, e.g., their smartphones. However, as it might be cumbersome to fill in an entire form using such small devices, De Luca and Frauendienst suggest to let users select which information they would prefer to fill in securely on their devices. The remaining data can be entered via the public interface of the display, e.g., the keyboard.

A different point of view is taken by Cao et al., as they are looking into ways of “enhancing privacy in public spaces through crossmodal displays” [43]. Their displays can be used by multiple users in parallel and support both public as well as personal information. The approach is rooted in the psychological concepts of *crossmodal* attention, i.e., people’s ability to correlate multiple sensory impressions. Their prototypical implementations, called *CrossFlow* and *CrossBoard*, use crossmodal cues to point out relevant information to individual users of public displays. *CrossBoard* highlights distinct pieces of information periodically. Each time the information that is particularly relevant to certain users is highlighted, their personal mobile devices vibrate. *CrossFlow* uses the same approach in the context of an ambient navigation system that projects a pattern of moving objects onto a surface. *CrossFlow* is thus comparable to the *Rotating Compass* as

presented by Rukzio et al. [189]. With regard to privacy, Cao et al. point out, that their system does not “need to sense or track the users, thereby maintaining user privacy” [43].

While scrutinizing the related work, it is also striking that users often do not regard privacy as an important topic. The following excerpts provide some evidence for this finding:

Our group appears to have a high tolerance for [perceived] privacy intrusion, since more than 90% of the group wears their badges regularly, and only one person has complained about the web cameras (and even that person appears to have grown used to them). —*McCarthy et al. [139]*

Overall, the number and range of comments about privacy issues was remarkably small. The presentation here might seem to make privacy a bigger issue than that which was expressed in the data. —*McDonald et al. [141]*

Only one participant raised security or privacy concerns. [...] Our initial concerns relating to privacy or user’s reluctance to change their names do not appear to have been founded. —*Davies et al. [63]*

Privacy is a factor that is mentioned often in the literature. It was thus somewhat unexpected that it was only mentioned by a single user in the study, and did not correlate strongly with display usage. [...] Even when asked users stated that privacy was less important. —*Müller et al. [153]*

The interviewees considered that, given the technology’s characteristics, privacy was a question of personal choice, and most weren’t concerned about it. —*José et al. [107]*

Finally, Nissenbaum also mentions similar “Puzzles and Paradoxes” in her book [161]. She points at the apparent discrepancy between people’s (theoretical) answers to questions about privacy and their (actual) reactions, habits, and actions. She also argues that this behavior is mainly due to a lack of knowledge or awareness:

One is a paradox, a stark contradiction at whose heart is this: people appear to want and value privacy, yet simultaneously appear not to value or want it. [...] What people do counts more than what they say, and what they do expresses quite the opposite of what is indicated by the polls. [...] On these grounds, computer scientist Calvin Gottlieb concludes, “I now believe that most of the populace really does not care all that much about privacy, although, when prompted, many voice privacy concerns.” [...] One reason for this is that people often are not fully aware that at certain critical junctures information is being gathered or recorded. Nor do they fully grasp the implications of the informational ecology in which they choose and act. Some claim that it is unfair to characterize a choice as deliberate when the alternative is not really viable; for instance, that life without a credit card, without a telephone, or without search engines requires an unreasonable sacrifice.

—*Nissenbaum [161, pp. 104–105]*

The review of related work presented in this section shows the breadth of research on threats and privacy issues related to public displays. However, there appears to be no universal threat model applicable to public display systems specifically, as none of the sources define or apply such a model. In her book, Nissenbaum mentions that “concern over privacy has also reached the scientific world of technical development and deployment, not only yielding a dedicated array of

privacy preserving technologies but also leading to the adoption of hardware and software design standards [...]” [161, p. 7]. This thesis contributes to this adoption of software design standards as described by Nissenbaum by proposing the privacy threat model (C1) as described in Section 10.1.

7.3. Personalization

In an attempt to classify and differentiate various approaches to personalization of public displays, Davies et al. [64] propose three personalization usage models. Table 7.1 presents each model in detail.

Walk-by and active personalization most often require means of interaction. The design space of this interaction spans two dimensions, as it can be either explicit or implicit, and direct or indirect: Explicit interaction includes, e.g., the active selection of contents by pressing buttons. Implicit interaction can be, for example, realized by analyzing the user’s posture, gaze, or other subconscious behavior. Direct interaction refers to situations in which the user gets in contact with the actual public display, e.g., by using its touchscreen. Finally, indirect interaction describes setups in which the user employs additional devices, e.g., smartphones, to personalize the public display. Each type of interaction, i.e., explicit, implicit, direct, or indirect, can be realized via various technical means. Common examples are mice, keyboards, touchscreens, speech or voice recognition, proxemics, and gestures.

Additionally, walk-by personalization also requires the public display to be able to identify individual users. Various technical means exist to facilitate this identification, for example, face recognition via cameras, *Bluetooth* or WiFi MAC address detection.

Table 7.1.: Personalization usage models by Davies et al. [64].

Usage Model	Definition
Walk-by	“Viewers passing by a single display see content that is relevant to them (as exemplified in the 2002 film <i>Minority Report</i> in which the characters are subject to personalised adverts as they journey across the city)” [64].
Longitudinal	“Users may express a preference to see personalised content and this is realised as an overall shift in the programming for a given geographic area such as a campus or shopping mall – typically over an extended period of time. The aim of such a usage model is not to try and ensure that any given display shows content for a specific user, but to try to ensure that, within a geographic region, the content viewers see is more representative of their interests than would be possible without personalisation” [64].
Active personalization	“Users (inter-)actively engage with a display system to control personalised applications on a nearby display, e.g. to extend a mobile phone display for better viewing of complex data” [64].

In contrast, longitudinal personalization does neither rely on identifying individual users nor on explicit or direct interaction. This approach rather uses sensors, such as cameras or proxemic sensors, that allow for detecting the overall interests, for example.

Rukzio et al. [190] present a matrix that correlates the number of persons that use a public display with the number of persons that may see the public display. Their matrix indicates what type of content—with respect to personalized information—would be appropriate to show.

For example, if a public display can be seen by a large number of persons, they suggest to present “personalized information that can be shown in public if no link to the initiator can be drawn [...]” [190]. In contrast to that, if only one person can see the public display, it would be acceptable to present “personalized information that must not be shown in public (e.g. automatic form filling of a form for ordering a book which shows address and bank account)” [190].

7.4. Design

In comparison to other traditional disciplines, public displays “represent a young and exciting area of research” [62]. Hence, researchers cannot draw on common or established methods, techniques, and tools in order to investigate research questions. Davies et al. [62] see three reasons for why researching public displays may be challenging:

No single goal. “Ads most likely strive to maximize public attention, interactive games may want to create an engaging experience, informative applications such as a public transport schedule may aim at maximizing usability and some displays may be designed to fade into the background, just presenting ambient information. Hence, metrics for display systems need to cope with different content, situations and purposes if meaningful comparisons are to be made” [62].

No models or simulations. People’s reactions to public displays are often of special interests to researchers and designers. In non-trivial situations, however, such reactions cannot be described in theoretical models since the behavior of individual users is very hard to predict. Therefore, “display systems often need to

be evaluated in the wild because there are no models or simulations that can be used for experimentation” [62].

Extremely challenging field studies. To evaluate public displays “in the wild” requires many efforts: robust prototypes have to be built, ethical regulations have to be obeyed, and consent of all stakeholders, see Section 2.2.1, has to be reached, for example. This usually requires a substantial amount of financial investment.

Despite these difficulties, researchers analyzed public displays in many respects. The remainder of this section focuses on common design considerations, design spaces, and design analogies.

With regard to design considerations, Huang et al. [98] found that research on public displays comprises (i) private use, (ii) semi-public use, and (iii) public use. Similarly, Dix and Sas [73] identified the importance of the spatial context of a public display and introduced these three degrees of *publicness*: (i) fully public, (ii) semi-public, (iii) semi-private. Systems that fall into the first category, are completely accessible for everyone, e.g., outdoor installations. Systems of the second category are only accessible to a certain group of people, e.g., employees or visitors of a company. The third category comprises systems that are supposed to be accessible to a very limited group of users, e.g., authorized medical staff. In the same paper, Dix and Sas use a similar threefold distinction to locate their work within a design space spanned by the dimensions “input device” and “display possibilities:” (i) personal, (ii) group, and (iii) public.

Similar to Dix and Sas, O’Neill et al. [168] propel the threefold notion of *information spheres*: (i) private spheres, (ii) social spheres, as well as (iii) public spheres. They “describe how information and services may be classified according to the kinds of access to them that are required” [168]. The private sphere contains “completely private issues,

services and information, access to which must be tightly controlled” [168]. The social sphere “contain[s] information relevant to a group of people. Social dynamics and constraints prohibit information in a social sphere from being made completely public” [168]. O’Neill et al. refrain from a very precise definition of public spheres. They merely claim that “the public sphere is not simply a collection of social spheres” [168]. Yet, they exemplify that “[...] a task such as looking up the train timetables would be included in the public sphere” [168].

A very comparable approach is taken by Azad et al. [17]. They analyzed the applicability of Hall’s proximity zones [88] to public displays. The results of their study indicate, that all four zones, i.e., the (i) intimate zone, the (ii) personal zone, the (iii) social zone, and the (iv) public zone, may also appropriately describe *territoriality* behaviors in front of public displays. They note, however, that “these zones must be modified to accommodate interactions around and on a large display given the frequency and acceptability of people violating each other’s personal and intimate zones” [17]. Such modifications mainly refer to an adjustment of the size of each zone.

Research on public displays frequently employs a public/private dichotomy, having a threefold or sometimes even finer level of granularity. This observation runs afoul of the conclusion presented in Section 6.2. However, it is beyond the scope of this thesis to evaluate whether the use of a public/private notion would be appropriate for public displays. Yet, as O’Neill et al. conclude, “a successful system [...] would necessarily reflect such aspects [‘the flexibility of real-life private/public distinctions’] of people’s everyday behaviour” [168].

Another strand of research scrutinizes the audience and its activities in front of public displays. Kaviani et al. [112] partition the audience into three groups: (i) actors, (ii) spectators, and (iii) bystanders. Members of the first group “feel encouraged by the display environment to take

an active role in the content. Actors may control and/or manipulate the content on these displays, e.g. by means of a hand held device, and so they can change the ‘flow’ and ‘pace’ of the presented content over time” [112]. People belonging to the second group “are mentally engaged with the displayed content and surrounding environment, but are not actively manipulating the content on the display” [112]. Finally, “bystanders are individuals who have no strong interest in the presented content on the display installation” [112].

Brignull and Rogers [38] introduce three kinds of activities around public displays: (i) peripheral awareness activities, (ii) focal awareness activities, and (iii) direct interaction activities. People that can be assigned to the first category are usually engaged in other (social) activities that take part in the vicinity of the public display and they do not pay attention to the system, cf. the group of bystanders introduced above. Members of the second group “are engaging in socializing activities associated with the display—talking about, gesturing to and watching the screen being used” [38]. This group resembles the spectators introduced above. Finally, the third group contains actors, i.e., people that interact with the public display.

Beyer et al. [27] looked at the audience behavior around cylindrical screens. Their work implies that many of the established assumptions for flat systems do not hold for differently shaped public displays. Regarding privacy, the fifth design assumption they present is especially interesting: “The position centrally in front of the display is preferred” [27] by actors. Consequently, content at the outskirts of the display may be perceived easily by spectators or bystanders.

The next part of this section focuses on design spaces in the domain of public display systems. Müller et al. [151] propose a design space for interactive public displays. It comprises two aspects: (i) mental models and (ii) interaction modalities. The mental models describe how peo-

ple perceive displays in a particular application scenario. Müller et al. report on these four common mental models: (a) posters, (b) windows, (c) mirrors, and (d) overlays.

They describe a poster as “a piece of printed paper [...], which can be attached to walls or vertical surfaces. Though electronic posters allow for a more dynamic content, they often show a mere adaptation of content created for their analog counterparts” [151]. Accordingly, windows provide “the illusion of a link to a remote, often virtual, location. In contrast to the poster, windows may work in two ways: users look inside, but windows offer the chance for the remote side to look outside as well” [151].

The third model, mirrors, pursues “the metaphor of a mirror to encourage interaction” [151] as it has been proven that “making users a part of the display has a strong potential to catch a user’s attention as they pass by” [151]. Finally, overlays “are frameless in that they can seamlessly integrate with the environment” [151], for example, by using projectors rather than computer displays.

The interaction modalities presented by Müller et al. are (1) presence, (2) body position, (3) body posture, (4) facial expression, (5) gaze, (6) speech, (7) gesture, (8) remote control, (9) keys, and (10) touch. As Müller et al. point out, these modalities are strongly related to the set of available sensors, e.g., touch, RFID, cameras, or microphones, and they are thus likely to change quickly as new technologies emerge. Eventually, they propose a taxonomy of public displays by extending their design space with another dimension: “type of supported interaction,” i.e., explicit or implicit interaction, see Section 7.3.

A particular aspect of a design space for public displays is concerned with the phases of user interaction and possible transitions between the phases. They are, for example, referred to as *interaction phases*

[151] or *interaction models* [5]. Sometimes, these interaction phases correlate to previously introduced design considerations. For example, Brignull and Rogers introduced three activity spaces that can also be found in their “model of public interaction flow” [38], as presented above. The *audience funnel*, as described by Michelis and Müller [146], is very frequently used in research and design processes. Figure 7.1 visualizes the audience funnel along with the *honeypot effect*, as introduced below. The audience funnel characterizes the different phases of interaction between public displays and their users. Depending on the specific public display system at hand, each phase bears particular challenges with respect to increasing the number of interactions. Their framework helps designers and researchers to identify the phase in which most users apt to abandon interaction. Furthermore, the audience funnel also provides a metric, that allows for a quantitative comparison of different public displays.

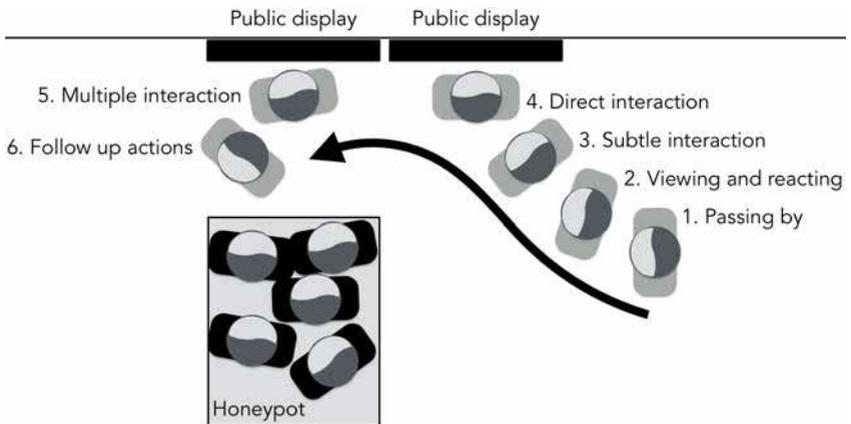


Figure 7.1.: Visualization of the audience funnel [146] (indicated by the arrow) and the honeypot effect. Adapted from [151].

As indicated by Figure 7.1, research observed a phenomenon called the honeypot effect [38, 225] while analyzing people’s behavior in front of and around public displays. Ten Koppel et al. describe this effect as

[...] The social effect of people being attracted to the public display by other people standing in close vicinity to it. It creates a social atmosphere around the public display in which people do not only signal their interest towards the display but also express that they are open for social interaction. Several field studies on public displays [24, 26] reveal that the honeypot effect is powerful in attracting users: once there is an initial crowd, people will be attracted by it and thereby again attract others [16]. —*Ten Koppel et al. [225]*

Thus, the honeypot effect ostensibly impacts the privacy of public display users, as other people may involuntarily tend to cast a glance towards the display and the shown contents. Ten Koppel et al. also discovered that the physical setup or constellation of public displays impacts the magnitude of the honeypot effect. This should, in turn, be considered when designing privacy-preserving public displays, see Section 9.2. Further studies also indicated that users usually pile up behind actors who are already engaged with the public display [17, 134]. This piling may raise privacy concerns, as actors—possibly handling sensitive data—will most likely not be able to sense the arriving spectators or bystanders without additional means, cf., e.g., the approach presented by Brudy et al. [40] in Section 6.2.

Table 7.2.: Public display user values by Perry and O’Hara [183].
ID prefix “AS-” omitted for brevity.

ID	Category	Description
A	Access to information	“[The] quick retrieval and [...] incidental access [to information]” [183] as well as the ability to “place information in contextually appropriate locations associated with particular tasks” [183].
A1	Retrospective reminding	“This refers to the cuing of memories of past events, places, time periods, activities and people, etc.” [183].
A2	Prospective reminding	“This refers to information displayed to remind the person to do something at a future point in time or when triggered by a particular context” [183].
A3	Display for take away	“Displayed information for taking away and use elsewhere. [...] For common spaces, it was common for printouts and contact cards to be left out on a reception area table for visitors to take away with them” [183].
A4	Quick reference	“A common function of displayed information was for referencing frequently required or difficult to remember information [...]. We commonly observed displayed lists of phone numbers [...]. Similarly, calendars and diaries of events were also displayed” [183].
A5	Learning	“Related to displaying information for quick reference is the notion of displaying things to learn. [...] As an example of this, one of the participants had pinned up her new fax number in an easy to see location to help her learn it as she referred to it” [183].

Table 7.2.: Public display user values by Perry and O’Hara [183] (continued). ID prefix “AS-” omitted for brevity.

ID	Category	Description
B	Social orientation	“[...] We can deliberately manipulate what our spaces visibly contain to present a social image of ourselves for a variety of purposes” [183].
B1	Identity and image	“Displayed information provides important information about identity of the owner [...]. The participants were very aware of what their displays expressed about themselves and what others might think [...]. Image management was an important concern as to what was appropriate and inappropriate to display” [183].
B2	Social grooming	“Another reason for displaying information was for social grooming and motivational purposes [...]. As well as the value of the display to it [sic] owner there were some interesting issues associated with visibly demonstrating appreciation for them” [183].
B3	Demonstrating achievements	“There were several examples of individuals, groups, and organisations displaying things to demonstrate their achievements. Certificates, exhibits, excellence awards and patent awards were observed in this kind of display behaviour” [183].

Table 7.2.: Public display user values by Perry and O’Hara [183] (continued). ID prefix “AS-” omitted for brevity.

ID	Category	Description
C	Co-ordination and planning	“[Displays] also have a role in co-ordinating the actions of different people within an organization. The key to understanding this is how displays make information about their creator’s current activities available to others so they can align their own activities with them accordingly. This may occur informally [...] or at an organisational level with formal status” [183].
C1	Communication and awareness	“Information is sometimes displayed [...] for communicating new information to others [...]. For example, ‘working at home-contactable on 555 75654’ [, ...] in-/out displays showed where people were, their holiday schedules or more general awareness information” [183].
C2	Conversational resources	“[...] A resource for initiating and scaffolding conversation. Visible information can be seen by visitors and invites comment [...]” [183].
C3	Current and past work processes	“[...] Displayed information associated with current and past work, including printouts, presentation slides, design sketches, research results, and whiteboard diagrams” [183].
C4	Planning and information overview	“[Displays] used specifically for providing an overview of certain information or activities. Project timelines and schedules were typical instances of this, as were ‘to-do’ lists” [183].

This section concludes by presenting a “taxonomy of visual display-based activity in office spaces” as proposed by Perry and O’Hara [183]. In their paper, the authors identified and evaluated analogue application scenarios (or user values as they refer to them) in office environments. Their findings can likely be applied, however, to digital application scenarios in other environments as well since the authors strive to derive “implications for the design of digitally enhanced and networked display technologies [...] from the findings [...]” [183]. Furthermore, Mark Weiser’s [237] vision of ubiquitous computing, i.e., computing devices that seamlessly blend into our everyday lives and replace their analogue counterparts, support this presumption. Moreover, their application scenarios, as presented in Table 7.2, may be used to cluster existing research projects. Section 7.5 pursues this approach to present related work on threats, threat models, and countermeasures. The taxonomy by Perry and O’Hara proposes these three key application scenarios for displays: (AS-A) *ready access to information*, (AS-B) *social orientation*, and (AS-C) *co-ordination and planning*.

7.5. Threats, Threat Models, and Countermeasures

As Section 6.5 indicated, a comprehensive analysis of threats, threat models, and countermeasures for a particular (computer) system may be painstaking. Public displays are usually complex systems, that are defined by numerous aspects, such as the given application scenario or the type of personalization. In general, the context of such systems significantly affects the list of impending threats and the set of suitable countermeasures. Thus—despite the heading of this section—, it would be too short-sighted to merely focus on threats, threat models, and countermeasures only.

7.5.1. Literature Survey

To provide a holistic view on this topic, this thesis presents the results of an extensive literature survey, that comprises 120 publications between the years 2000 and 2014. The results of the survey provide new insights into two domains: (i) the distribution of application scenarios and the types of personalization being applied; and (ii) the distribution of privacy threats as well as proposed countermeasures. In both domains, this section points out areas that have been well covered by research as well as those that require further work. For example, privacy threats that have not been evaluated in a particular application scenario or that have not been addressed by any countermeasure may be of special interest to the research community, as they may point to research opportunities. Also, designers may use the findings as a reference when designing, prototyping, or evaluating public display systems. This section may thus serve as a guide for researchers as well as for designers of public displays, both striving for privacy-preserving personalized systems.

Up to this point, this thesis illustrated the breadth of research on personalized public displays. This section surveys recent related work by applying a novel classification scheme. The scheme is based on two established categorization dimensions. The first dimension spans the user values, i.e., the twelve application scenarios as defined by Perry and O'Hara [183], presented in Table 7.2. Clearly, it is impossible to draw a sharp line between each of the twelve application scenarios, as a surveyed project could be mapped to multiple applications or it rather fits "in-between." However, the objective of this section is not to establish a normative definition of applications for personalized public displays. It rather uses the application scenarios suggested by Perry and O'Hara to cluster existing research projects to support the two main goals of this section: (i) to provide an overview of re-

search on personalized public display systems with a special focus on privacy issues and potential countermeasures; and (ii) to identify areas for further work with regard to privacy threats in order to design privacy-preserving personalized public display systems.

The second dimension of the classification scheme comprises the personalization usage models as introduced by Davies et al. [64], see Table 7.1. This dimension may help to determine whether public displays of a certain personalization type are more or less likely to be subject to certain privacy threats. A possible result of the analysis could be, for example, that public displays that fall into the category of active personalization are more prone to privacy threats than longitudinal systems, as the latter ones may offer fewer publicly available interfaces, such as keyboards or mice.

While early installations of public displays occurred in the late 1970s [195], this survey focuses on publications and work between the years 2000 and 2014. The key reason for this is the pace of technological development: Besides form factors, like screen sizes and shapes, there has been a significant progress in technology, especially with the advent or proliferation of the Internet. Additionally, various means of interaction appeared in that time frame, e.g., smartphones or depth cameras, such as Microsoft Kinect. As explained in Section 7.3, interaction is an important means to actually personalize systems. As a final remark, the survey does not include work related to media facades, as these instances of public displays are mostly used for artistic installations and are usually not designed for personalization.

7.5.2. Visualization of the Surveyed Work

This survey analyzed 120 papers on public displays published between 2000 and 2014. Then, 68 out of 120 papers were selected for review, focusing on those, which investigated personalized public display systems. Table 7.3 provides an overview of the results, i.e., all projects, application scenarios, personalization usage models (P-Type), privacy threats (see Section 6.5), and countermeasures. For brevity, countermeasures are referenced by numbers. The corresponding names are listed in Table 7.10.

The hive plot shown in Figure 7.2 supports the reader by (i) establishing an overview of the resulting distribution in terms of application scenarios (blue axis, 12 o'clock), privacy threats (green axis, 4 o'clock), and countermeasures (gray axis, 8 o'clock), and by (ii) highlighting the most significant relations between individual items on the axes. The size of an item on each axis corresponds to the number of papers that fall into that category. For example, there are twelve papers in AS-C4 and six papers in AS-B3. Thus, the area of the bubble labeled AS-C4 is twice as big as the area of bubble AS-B3. The stroke width of each link is directly related to the number of corresponding papers. For example, the link between AS-A4 and I is almost twice as thick as the link from AS-A5 to I (representing 20 vs. 11 papers). To reduce the visual clutter, links representing less than seven papers are grayed out, where seven is the upper median of all link widths. The upper median was chosen as it is more robust against discordant values and it represents an actual value contained in the underlying set of figures. As a consequence, Figure 7.2 highlights links whose significance is "above average." The median value of seven was also used to demarcate the items with highest counts in the Tables 7.4, 7.8, and 7.9.

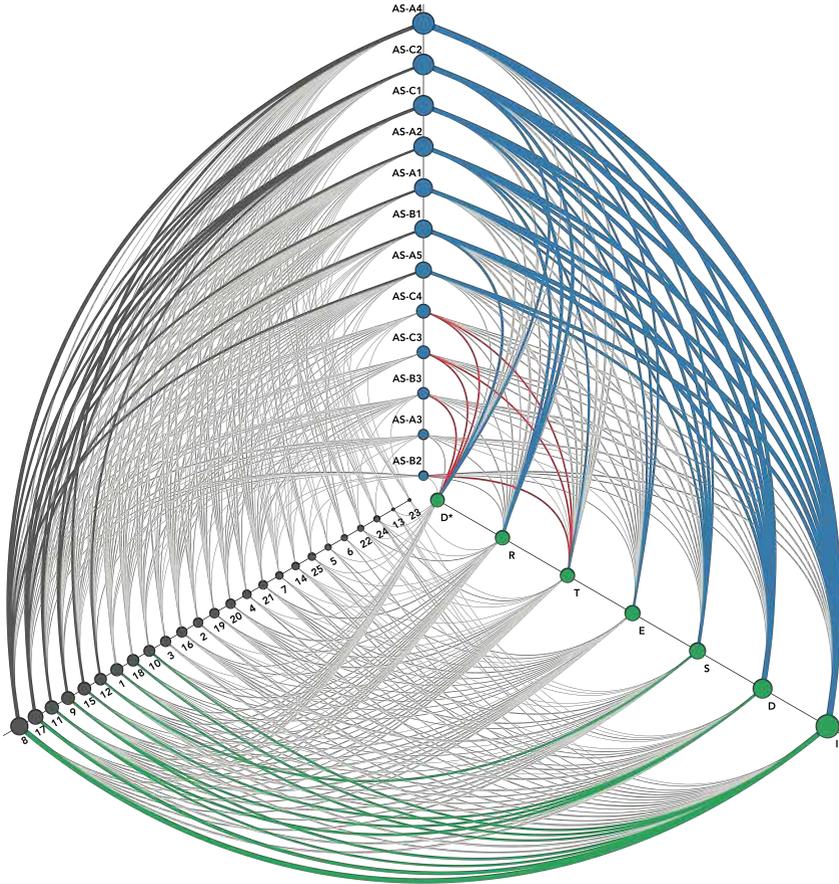


Figure 7.2.: Hive plot visualizing the surveyed scientific work with regard to application scenarios, threats, and countermeasures. Red lines indicate research opportunities.

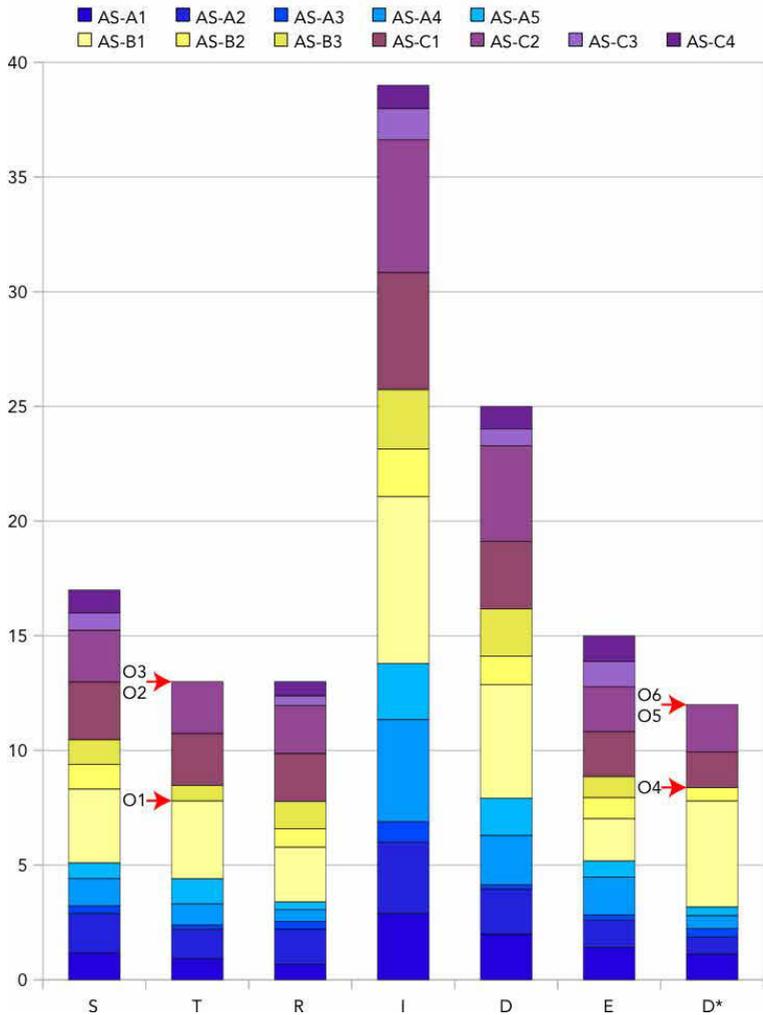


Figure 7.3.: Distribution of application scenarios and threat categories. Red arrows indicate six research opportunities O1–O6.

Figure 7.3 visualizes the harmonized distribution of each application scenario to the privacy threats introduced by the STRIDED* model. Total amounts of publications have been harmonized across major application scenarios (i.e., A, B, and C, see Table 7.2) to compensate for mixed sample sizes. On the ordinate (0–39), each column thus represents the total number of publications that can be assigned to a STRIDED* category; values indicated by individual column segments, however, do not represent absolute values to be directly related to Table 7.3, but relative values that allow for comparing publication counts. Red arrows indicate the six apparent research opportunities as discovered in this survey: (O1) tampering (T) with respect to social grooming (AS-B2); (O2) tampering (T) with respect to current and past working processes (AS-C3); (O3) tampering (T) with respect to planning and information overview (AS-C4); (O4) decontextualization (D*) with respect to demonstrating achievements (AS-B3); (O5) decontextualization (D*) with respect to current and past working processes (AS-C3); (O6) decontextualization (D*) with respect to planning and information overview (AS-C4). Figure 7.3 will be explained in more depth when presenting and discussing the results on application scenarios and threats.

7.5.3. Outcomes and Discussion

This subsection presents and discusses the outcomes of the survey with respect to (i) application scenarios, (ii) personalization usage models, (iii) threats, (iv) countermeasures, and (v) interrelations. The structure of each segment is derived from the structure implied by Figure 7.2: Items emphasized by color are addressed first, followed by the remaining (uncolored) items. Finally, this subsection highlights research opportunities, also visualized in Figures 7.2 and 7.3.

Application Scenarios

Outcomes. The analysis of the surveyed work reveals which of the twelve application scenarios is mentioned most often, see Table 7.4. The numbers of publications in each application scenario do not add up to the total number of surveyed papers (68 publications were examined), since each paper may be assigned to more than one application scenario. For example, *Digifieds* [7] falls into the two categories AS-A2 (prospective reminding) and AS-A3 (display for take away).

When analyzing the harmonized distribution of each application scenario over the privacy threats introduced by the STRIDED* model, a uniform distribution can be observed, see Figure 7.3. By comparing the heights, i.e., the amount of publications, of each application scenario per threat type, there appear equal amounts of publications per threat. This is an example for the threat of spoofing, i.e., the first column in Figure 7.3: There are 6 publications mentioning this threat type in all application scenarios related to access to information (indicated by five blue segments). There are also six publications discussing this threat type in all application scenarios related to social orientation (indicated by three yellow segments). Finally, there are seven publications mentioning this threat type in all application scenarios about co-ordination and planning (indicated by four purple segments).

Discussion. First of all, this dimension of the survey classification scheme seems to suit the body of related work well, i.e., every public display system referred to in a particular publication can be assigned to either of the twelve categories. Yet, the total amount of publications addressing each of the three major application scenarios varies significantly. The application scenario access to information (AS-A) has been applied 55 times, while social orientation (AS-B) has been

Table 7.4.: Application scenarios ordered by their frequency of occurrence in literature. A rule separates the most frequent application scenarios from the remaining ones.

ID	Description	Count
AS-A4	Quick reference	33
AS-C2	Conversational resources	30
AS-C1	Communication and awareness	28
AS-A2	Prospective reminding	26
AS-A1	Retrospective reminding	22
AS-B1	Identity and image	21
AS-A5	Learning	17
AS-C4	Planning and information overview	12
AS-C3	Current and past work processes	10
AS-B3	Demonstrating achievements	8
AS-A3	Display for take away	6
AS-B2	Social grooming	5

assigned 25 times, and co-ordination and planning (AS-C) has been addressed 45 times. These aggregated figures are not shown in Table 7.4. The sum of all totals exceeds the number of surveyed papers, since papers may be assigned to more than one application scenario.

In particular, 33 publications are related to the application scenario quick reference (AS-A4), and 30 publications analyze the use of conversational resources (AS-C2). In contrast, social grooming (AS-B2), display for take away (AS-A3), and demonstrating achievements (AS-B3) appear to be the three least covered application scenarios. This indicates a special interest in public displays as a means for information retrieval (AS-A) and collaborative work (AS-C1, AS-C2). Applications

scenarios with aspects of social orientation (AS-B) seem to spark less interest. Accordingly, the “most prominent” application scenarios are highlighted in Table 7.4.

This finding goes along with the intuitive observation that most public displays are nowadays used to disseminate information to the public, for example, flight departure times at airports or news feeds in subway stations, and to let people jointly work on a specific task, e.g., on a mind map. In the light of ever growing numbers of social network users, it is an interesting question why there is relatively little research on the application scenario of social orientation on public displays. Despite all privacy concerns, people tend to disclose sensitive information on those platforms. It might thus be interesting to conduct research in this area to see how personalized public displays may influence social environments.

Harmonizing the varying sample sizes of each application scenario reveals insightful results as well. Most privacy threats have been analyzed with the same intensity in each application scenario. Ideally, 33 percent of each column in Figure 7.3 should be designated to one of the three major application scenarios—which roughly is the case, see percentages provided on p. 114. However, there are three exceptions that may be potential fields for further research: elevation of privilege in AS-B (24.54 percent), repudiation in AS-A (26.10 percent), and decontextualization in AS-A (26.45 percent). (These percentages are not depicted in Figure 7.3.) In addition to these slight imbalances, there are also research opportunities relating to combinations of application scenarios and privacy threats that have not been covered yet. These research opportunities will be discussed below in more detail.

According to the results presented in Subsection 10.1.1, users rank privacy as most important for actions (application scenarios) that involve personal messaging, browsing pictures, social networking, calendar-

ing, and watching videos. Compared to this ranking, the prioritization based on the results of the survey highlight a number of matches and mismatches, as shown in the Tables 7.5 and 7.6. The mapping of calendaring to quick reference, for example, is based on a keyword analysis of the original definition proposed by Perry and O’Hara [183]. All of the actions with presumably high demands for privacy (according to the prioritization presented in Section 10.1) were covered—at least partially—by the surveyed research. Yet, some aspects of certain actions, e.g., the aspect of social grooming while using social networks, were not covered so far. Any of the combinations listed as mismatches thus constitute unexplored research opportunities.

In the opposite direction, analyzing the ranking (Section 10.1) in the light of the results of this survey also yields some interesting findings. Some of the proposed actions, such as browsing websites, checking itineraries, getting directions, and using a map, can be mapped to the application scenario quick reference (AS-A4). Though users do not seem to require a high level of privacy for either of these actions, they have been well covered by research. Future work may thus focus on the remaining mismatches listed in Table 7.6 first.

Table 7.5.: Matches between application scenarios analyzed in the literature survey and actions with presumably high demands for privacy.

Application scenario	Action
Quick reference (AS-A4)	Calendaring
Communication and awareness (AS-C1)	Personal messaging
Retrospective reminding (AS-A1) Prospective reminding (AS-A2) Identity and image (AS-B1) Conversational resources (AS-C2)	Browsing pictures, watching videos
Retrospective reminding (AS-A1) Prospective reminding (AS-A2) Identity and image (AS-B1) Communication and awareness (AS-C1) Conversational resources (AS-C2)	Social networking

Table 7.6.: Mismatches between application scenarios analyzed in the literature survey and actions with presumably high demands for privacy.

Application scenario	Action
Social grooming (AS-B2)	Browsing pictures, watching videos
Social grooming (AS-B2) Demonstrating achievements (AS-B3)	Social networking

Personalization Usage Models

Outcomes. The survey indicates that the majority of public display systems reported on in literature rely on active personalization, see Table 7.7. Walk-by is the second most often used personalization usage model, followed by systems employing the longitudinal approach.

Discussion. This dimension of the survey classification scheme also fits well to the body of related work, i.e., every public display system referred to in a particular publication can be assigned to either of the three models. Table 7.7 indicates that the large majority of public displays apparently uses active personalization (A). This seems natural, as many of the systems deployed in the real world commonly use interfaces such as touchscreens, keyboards, or mice. These means of interaction cater for active personalization, for example, by selecting preferred contents via the click of a mouse. In contrast, sensors required for longitudinal (L) or walk-by (W) personalization are less common, e.g., RFID, Bluetooth, or depth cameras. Future research could thus shift its focus from active to longitudinal or walk-by personalization on public display systems.

Table 7.7.: Personalization usage models by frequency of occurrence.

ID	Description	Count
A	Active personalization	59
W	Walk-by	26
L	Longitudinal	6

Threats

Outcomes. The privacy threats mentioned in literature are summarized in Table 7.8. The threat of information disclosure has been analyzed most often, followed by the threats of denial of service and spoofing. A general observation of the survey is that earlier publications tend to focus more on the general feasibility and the technical boundaries than on security or privacy. A large body of research is concerned with evaluating the general use of public displays in various application scenarios [87, 139, 205], while another strand of research analyzes the users' reactions to this technology, see, e.g., first publications on *Instant Places* [107]. Subsequently, more nuanced analyses in terms of security and privacy emerged in subsequent years, see, e.g., follow-up publications of *Instant Places* [108].

Discussion. Figure 7.3 shows that a significant amount of research focused on the threats of information disclosure (39 publications), denial of service (25 publications), and spoofing (17 publications). These “most imminent” privacy threats are listed at the top of Table 7.8. The findings generally align well with the prioritization of threats as proposed in Section 10.1: The results suggest to primarily focus on threats induced by information disclosure and spoofing. This congruence indicates that research has tackled the privacy threats users fear the most on personalized public display systems.

However, there is also a partial mismatch between the prioritization presented in Section 10.1 and the findings of the survey. According to the prioritization, denial of service has been rated second to last in terms of privacy threats when designing personalized public display systems. This survey revealed, however, that a significant amount of research has been spent on exactly this privacy threat. Future research

could thus focus more on the remaining privacy threats according to the prioritization in Section 10.1, i.e., repudiation (R, 13 publications up to now) and decontextualization (D*, 12 publications up to now).

Table 7.8.: Privacy threats ordered by their frequency of occurrence in literature. A rule separates the most frequent privacy threats from the remaining ones.

ID	Description	Count
I	Information disclosure	39
D	Denial of service	25
S	Spoofing	17
E	Elevation of privilege	15
T	Tampering	13
R	Repudiation	13
D*	Decontextualization	12

Countermeasures

Outcomes. The survey also provided insights into countermeasures, that may be applied to counter some privacy threats. Table 7.9 lists all of the surveyed countermeasures. The following paragraphs will explain each countermeasure in more detail. A very common way to alleviate privacy risks for personalized public displays is to avoid showing personalized content on the public display itself by using a 2nd device (*No. 8*) instead. Examples of such a second device are smartphones [43] and wearable devices [26]. Restricting access (*No. 17*) to public displays to certain locations (e.g., locked rooms [50]), groups of users (e.g., family members or employees [138, 163]), or individual users [139] is the second most often applied countermeasure.

Another very common approach is to let social protocols, conventions, or norms (*No. 11*) take care of possible privacy threats [180]. For example, social norms and good manners may imply that one should exercise discretion, if a public display, e.g., an interactive kiosk, is already in use by another person. The terminal itself, however, does not provide any means to account for such a situation. In case that public displays allow users to post personal messages, e.g., texts or images, such postings are very often moderated (*No. 9*) by staff [9]. Moderators may quickly decide whether certain contents are inappropriate (also based on context, which may be complicated for a machine or algorithm) and delete it instantaneously.

Sometimes, privacy and security issues arise because of an inappropriate design applied in early stages of the system development. One specific design flaw for personalized public displays is an interaction concept intended for single users only. Thus, applying an explicit multi-user design (*No. 15*) as early as possible is often recommended or applied [165]. A further common countermeasure is to employ

trusted, external services (*No. 12*), e.g., Facebook, Twitter, or Google, to take care of vital processes of a personalized public display, for example, registering users or retrieving real names [97].

To do nothing (*No. 1*) or to refrain from any countermeasure on purpose appears to be another approach often used [22, 215]. Countering privacy issues with abstract representations of information (*No. 18*), e.g., via color codes, regular patterns, or figurative metaphors, is another means [60, 189]. Anonymizing users or their data (*No. 10*) may also help to mitigate privacy risks. For example, users could contact each other by box-numbers rather than phone numbers or e-mail addresses [7], or their identity would be indicated via silhouettes rather than actual photographs [123].

The next countermeasure on the list corresponds to masking (*No. 3*) sensitive content, for example, by blacking out specific screen regions [40] or by using visual filters to completely block its perception [205]. Similar to countermeasure No. 15 is the approach to apply a minimal data gathering design (*No. 16*) at early stages of the development process. This way, the personalized public display tries to avoid unnecessary data aggregations and offers its service while processing as little data as possible [38, 64, 217].

Closely related to masking (*No. 3*) sensitive content is the approach to minimize (*No. 2*) it. The difference is, that the minimized content is no longer visible to anybody (while masked content would still be perceivable from close distance), but its presence can still be detected, for example, indicated by an icon [40, 185]. One common way to mitigate privacy threats induced by tampering with hardware is to use protective casing (*No. 19*) [50, 223].

The next countermeasure is very common as it is not limited to public display systems, but can be applied to any kind of computing system:

logging or auditing (*No. 20*). Systems equipped with this countermeasure meticulously record all kinds of events, for example, interaction times or user input. This way, it might be possible to detect or monitor a privacy breach and trace back its cause as well as to detect actual user behavior in order to adapt accordingly [6, 50, 152].

To completely blind out personal data (*No. 4*) instead of masking or minimizing it in threatening situations may be another means [40, 185]. To have users acknowledge (*No. 21*) certain contents may be a countermeasure applicable to some privacy issues, such as repudiation [50, 122, 152]. If users are able to explicitly assign, claim, or “carve off” (*No. 7*) screen real estate to use it while interacting with a personalized public display, this may help to remedy possible privacy risks such as denial of service or elevation of privilege [103, 231].

Another countermeasure reported on in literature is to require explicit UI interactions (*No. 14*). For example, after deleting a personal text message on a public display, the system should not automatically show the next message, as it is a common behavior of desktop e-mail clients. Instead, it should wait for the user to explicitly request the next message [49, 50]. In some situations, the users’ privacy may be at risk, if the interaction itself can be observed by others, for example, since using the system may be embarrassing [123, 205]. In those situations, unobtrusive interaction (*No. 25*) may be a useable countermeasure.

Closely related to the countermeasure of removing all personal data (*No. 4*) is to blind out all data (*No. 5*) [121]. The next countermeasure proposes to raise the awareness (*No. 6*) for privacy threats, for example, by indicating that someone else is looking at the public display [40, 185]. In order to alleviate privacy threats induced by defamation etc., it may be helpful to de-anonymize users (*No. 22*) [230].

Table 7.9.: Countermeasures ordered by their frequency of occurrence in literature. A rule separates the most frequent countermeasures from the remaining ones.

ID	Description	Count
8	2 nd device	20
17	Restrict locations/groups/users	14
11	Let social protocols handle it	11
9	Moderation	10
15	Explicit multi-user design	10
12	Web-Of-Trust	8
1	Do nothing	9
18	Abstract presentation	8
10	Anonymize user data	7
3	Mask	6
16	Minimal data gathering design	6
2	Minimize	5
19	Protective casing	5
20	Logging/Auditing	5
4	Blind out personal data	4
21	Acknowledging	4
7	Assign/Claim/Carve	3
14	Require explicit UI interactions	3
25	Unobtrusive interaction	3
5	Blind out all data	2
6	Raise awareness	2
22	De-anonymize users	2
24	Human supervision	2
13	Plausible deniability	1
23	Restrict interaction	1

Table 7.10.: Heat map visualizing the use of countermeasures in the surveyed work to address privacy threats identified by the STRIDED* model.

ID	Description	S	T	R	I	D	E	D*
1	Do nothing	2	2	3	7	3	3	3
2	Minimize	0	1	0	5	1	2	1
3	Mask	1	1	0	6	0	2	1
4	Blind out personal data	0	1	0	4	0	2	1
5	Blind out all data	0	0	0	2	0	2	1
6	Raise awareness	0	1	0	2	0	0	0
7	Assign/Claim/Carve	0	0	1	2	2	2	0
8	2 nd Device	5	4	2	16	6	4	6
9	Moderation	6	3	5	8	7	3	5
10	Anonymize user data	2	2	3	7	4	0	1
11	Let social protocols handle it	5	4	5	7	8	6	4
12	Web-Of-Trust	7	4	3	6	5	3	4
13	Plausible deniability	1	1	0	1	1	0	1
14	Require explicit UI actions	1	1	1	3	2	1	0
15	Explicit multi-user design	3	2	1	5	7	6	2
16	Minimal data gathering design	5	2	2	3	5	2	2
17	Restrict locations/groups/users	3	2	4	11	6	3	2
18	Abstract presentation	1	2	2	7	3	1	0
19	Protective casing/restrict access	2	4	2	2	4	2	0
20	Logging/audit	3	2	3	2	2	2	0
21	Acknowledging	3	1	3	2	3	3	0
22	De-anonymize users	1	1	2	0	1	0	1
23	Restrict interaction	0	0	1	0	1	0	0
24	Human supervision	1	2	1	2	2	0	1
25	Unobtrusive interaction	1	0	1	3	1	2	1

In some situations, it might be appropriate to use human supervision (*No. 24*) in order to guarantee a proper and privacy-preserving use of a personalized public display [38, 223]. Next, it might be advisable to show other (unrelated) content besides the information actually requested by users. This way, users may employ the principle of plausible deniability (*No. 13*) to protect their privacy [64]. Finally, restricting interaction (*No. 23*) to a predefined set of actions may be used to counter certain privacy threats, such as denial of service [56].

Discussion. Based on the findings reported above, it is possible to derive a heat map of countermeasures, which is shown in Table 7.10. This heat map visualizes the correlation between each of the twenty-five countermeasures and all of the seven privacy threats. Counting the applications of one countermeasure to a privacy threat gives the magnitude of that particular combination, which in turn determines the applied color. For example, using a 2nd device (*No. 8*) has been applied 16 times to alleviate the privacy threat of information disclosure (I). Computing the magnitude for all other combination reveals that 16 actually is the largest figure and is thus assigned to the darkest shade of red. The magnitude of 0 is assigned to the color white. All remaining magnitudes in between are assigned to their corresponding colors in a linear fashion.

Researchers interested in personalized public display systems that respect privacy as well as designers of such systems can use this heat map to quickly gain insights and draw conclusions. For example, researchers may focus on the “white spots,” as they point to combinations of privacy threats and countermeasures that have not been researched yet—or futile combinations. Designers may use the heat map as a reference when designing, prototyping, or evaluating public display systems. Once they identified that their system may be

subject to a particular privacy threat, they can look for “the hottest countermeasure,” i.e., the one with the largest magnitude, within the corresponding column, i.e., either S, T, R, I, D, E, or D*, in Table 7.10.

The subset of the three most frequently applied countermeasures in literature consists of: (i) using a 2nd screen (No. 8), (ii) restricting access to certain locations, groups, or users (No. 17), and (iii) letting social protocols, norms, or conventions handle possible conflicts between demands for privacy and personalization (No. 11).

Though Davies et al. [62, p. 91] put a special focus on the first one (No. 8) and Perry and O’Hara [183] recommend the last one (No. 11), this particular subset of countermeasures seems to counteract the purpose or the intention of using public displays to some degree. To use another screen for showing sensitive content may indeed result in an increase of the users’ privacy, as the content is not shown in public anymore. However, one could argue that this increase has been achieved by simply avoiding the usage of the public display.

Based on the definition of a public display, restricting access to certain users, e.g., by locking it in a room [50], could render the public display a desktop computer screen, as it is not truly public anymore. Yet, as with the public/private dichotomy introduced in Section 6.2, demarcating a display as private or public may be a complex topic, see, e.g., literature by Huang and Mynatt [99] on *semi-public displays*. Lastly, though the application and evaluation of social protocols, norms, or conventions as a means to counter privacy threats on personalized public displays is an interesting strand of research [180], it may be less suitable for real world deployments. For instance, trusting in social norms may be a proper design for a shopping mall information kiosk, but it may be less appropriate when designing an ATM. All of the “most often used” countermeasures discussed above are listed at the top of Table 7.9.

Interrelations and Research Opportunities

Outcomes. A close look at Figure 7.2 reveals a complete bi-directional relationship between the colored items on each axis. The application scenario AS-A4 (quick reference) is linked to the privacy threats I (information disclosure), D (denial of service), and S (spoofing) on the one hand, and to the countermeasures No. 8, 17, 11, 9, 15, 12, 1, 18, and 10 (see Table 7.10) on the other hand. Similarly, the privacy threat I (information disclosure) is linked to all of the application scenarios AS-A4, AS-C2, AS-C1, AS-A2, AS-A1, AS-B1, and AS-A5 on the one side, as well as to the countermeasures No. 8, 17, 11, 9, 15, 12, 1, 18, and 10 on the other side. Finally, the countermeasure No. 8 (2nd device) is linked to the privacy threats I (information disclosure), D (denial of service), and S (spoofing), as well as to the application scenarios AS-A4, AS-C2, AS-C1, AS-A2, AS-A1, AS-B1, and AS-A5.

Thus far, this survey on personalized public display systems provided insights into immediately observable facts and interrelations, e.g., the total amount of papers concerned with a particular application scenario or the relationship between a privacy threat and applicable countermeasures. However, studying things that are less evident may also yield interesting results. As mentioned in the previous subsection—and as illustrated by red lines in Figure 7.2 as well as red arrows in Figure 7.3—, some application scenarios are not linked to particular privacy threats and vice versa. This signifies that none of the surveyed work is concerned with the specific combinations of application scenarios and privacy threats, as summarized in Table 7.11. Note that Figure 7.2 only visualizes these six research opportunities, i.e., the privacy threats that have not been addresses in a specific application scenario, in order to avoid visual clutter. Some weaker interrelations from Figure 7.2 were left out. Examples of such weaker interrelation would be “all countermeasures that have not been tested against a certain pri-

vacy threat” or “all countermeasures that have not been applied in a certain application scenario.”

When looking at maximum and minimum numbers (see Figure 7.3), it can be observed that there is less research on the privacy threat of repudiation (R) with regard to application scenarios classified as access to information (AS-A, 26 percent). Similarly, there are less publications investigating elevation of privilege (E) in the context of application scenarios classified as social orientation (AS-B, 25 percent). Finally, a comparatively small number of publications is concerned with the threat of decontextualization (D*) with respect to application scenarios related to access to information (AS-A, 26 percent). The percentages are not shown in Figure 7.3 for the sake of readability.

Table 7.11.: Research opportunities regarding privacy threats (T stands for tampering; D* stands for decontextualization) and application scenarios. The IDs correspond to the labeled arrows in Figure 7.3.

Privacy threat	Application scenario	ID
T	AS-B2 (social grooming)	O1
	AS-C3 (current and past working processes)	O2
	AS-C4 (planning and information overview)	O3
D*	AS-B3 (demonstrating achievements)	O4
	AS-C3 (current and past working processes)	O5
	AS-C4 (planning and information overview)	O6

Discussion. The complete bi-directional relationship between the colored items shown in Figure 7.2 indicates that previous research has independently addressed the privacy issues of most concern to users of public display systems:

1. All of the “most imminent” privacy threats have been analyzed in the “most prominent” application scenario AS-A4.
2. The “most imminent” privacy threat I has been analyzed in all of the “most prominent” application scenarios.
3. All of the “most often used” countermeasures can be used to alleviate the “most imminent” privacy threat I.
4. The “most often used” countermeasure No. 8 can be used to alleviate all of the “most imminent” privacy threats.
5. The “most often used” countermeasure No. 8 can also be applied to the “most prominent” application scenarios.
6. All of the “most often used” countermeasures can be applied in the “most prominent” application scenario AS-A4.

As mentioned above, the analysis of privacy threats and applications scenarios points to six research opportunities, labeled O1–O6 in Figure 7.3. The remainder of this section will explain each research opportunity in more detail.

O1. Third parties may tamper with processes of AS-B2 (social grooming) for a number of reasons. A very obvious motivation would be to want to influence the relationship between two people in a negative way, possibly in order to take advantage of somebody. The *Thank You Board* [157] is an example of such a social grooming application. Let

us assume that Alice has lent Bob her car. Bob would like to send Alice a grateful note via the Thank You Board. However, an attacker could have obscured the display of the Thank You Board, so that not all or even none of the notes can be seen. Alice may thus not get Bob's expression of thanks, which in turn may have a negative impact on their relationship, e.g., Alice may refuse to lend things to Bob in the future. It would thus be interesting to see in which ways tampering could be used to manipulate the social relations between users of personalized public displays. One interesting research question in this direction could be how robust social bonds are to interpersonal frictions induced by digital media.

O2/O3. The two application scenarios AS-C3 (current and past work processes) as well as AS-C4 (planning and information overview) both belong to the general category AS-C (co-ordination and planning). This general category is strongly related to the research field of CSCW. A common assumption in this research area is, that users of such systems are somehow related to another. For instance, they are employed by the same company or are co-workers on the same project. Based on this premise, the need for a tampering-resistant system may be less apparent compared to scenarios in which complete strangers share a public display. The *MERBoard* [100], for example, was designed to facilitate efficient collaboration of NASA employees working on a specific project. As the targeted group of users was limited and unlikely to sabotage the system, the project did not consider tampering as a threat. Following this argumentation, it may be acceptable to not analyze tampering in these application scenarios. However, one might also think of scenarios in which the premise stated above might not hold. For example, guests or visitors can gain access to a public display system, or frustrated employees can try to sabotage their

companies from within. An interesting strand of research would thus be to explore how personalized public displays could become resistant to tampering in application scenarios related to CSCW, for example, AS-C3 and AS-C4, or how intrusion (attempts) could be detected and communicated to the users.

O4. Depending on a specific situation, AS-B3 (demonstrating achievements on personalized public displays) can be a sensitive topic. For example, showing diplomas or PhD certificates to somebody may be regarded as appropriate if it can be assumed that both parties are business partners. In this case, the demonstration of achievements serves a professional purpose, e.g., highlighting a person's skills. Showing the same information may be regarded as inappropriate, however, if it could be assumed that both parties are friends and not business partners. In this case, the demonstration of achievements can be perceived as pretentious and inappropriate. Thus, an intriguing research questions would be, whether and how the context can be communicated to all users as well as onlookers or passersby of personalized public displays. In contrast, it would be interesting to analyze ways to properly shield personalized content from non-active users, e.g., onlookers or passersby, without a negative impact on the user experience of others. Such a negative impact is often seen when applying visual privacy filters to laptop screens as they often reduce the visual fidelity for users located right in front of the display as well.

O5/O6. As mentioned above, the application scenarios AS-C3 (current and past work processes) and AS-C4 (planning and information overview) both fall into the category of AS-C (co-ordination and planning), which is related to the field of CSCW. Moreover, in the context

of this domain, most users of personalized public displays are usually related to each other. Assuming a common ground of knowledge or intentions, it seems understandable that previous research has not focused on the privacy threat of decontextualization in these two particular application scenarios. Nevertheless, such situations can still occur. For example, two colleagues are using a public display to argue about the budget for a project. In the course of using the display one of them unwittingly brings up the budget plan of the previous project, e.g., by accidentally selecting the wrong file. That plan may provide significantly higher staff costs than the plan they are currently working on. Without the proper context, the other colleague may be led to think that his co-worker is now trying to compensate for these costs in the current project, which might not be the case. An interesting strand of research would be to evaluate how the threat of decontextualization could be alleviated in these application scenarios. In contrast to the situation outlined above for AS-B3, it may not be sufficient to simply shield particular pieces of information from the view of other users, as these users may be the intended receivers, but they do not know how to correctly interpret the information itself.

7.5.4. Limitations

Though this extensive survey of 120 papers on privacy and personalized public displays was executed with great care, it is still subject to some limitations. Though a large number of outlets was thoroughly scanned, it is possible that some relevant work of the last 15 years is not included. Additionally, the results might have been different, if publications prior to the year 2000 had been included. The focus was on work published after 2000 as prior publications often had a strong technical mindset, and did not consider privacy aspects at all.

The choice of another classification scheme, e.g., the “space of input device and display possibilities” as proposed by Dix and Sas [73], could have had an impact on the overall findings as well. In the same way, visualizing and analyzing more missing relations in Figure 7.2 could have yielded additional or other results. Finally, there is no statistical analysis of the data collected about the surveyed papers due to small sample sizes. Instead, there is a qualitative analysis, which led to the insights reported above.

7.5.5. Summary

As this survey of literature shows, privacy on personalized public displays has gained interest between the years 2000 and 2014. Early publications often focused on the technical feasibility and application scenarios rather than on security and privacy aspects. However, the awareness of privacy issues in the general public has increased in recent years, e.g., in the aftermath of disclosures on governmental surveillance in 2013. Privacy has thus become a sensitive topic regarding the design of public displays. To gain some insights into this vast research area, this section surveyed 68 publications on personalized public display systems by applying a novel categorization scheme. This scheme comprises the two dimensions of (i) application scenarios and (ii) personalization usage models, which proved to be useful to derive the findings of this survey.

This section also presented and discussed the most prominent application scenarios, the most common personalization usage models, the most imminent privacy threats, and the most often used countermeasures. The section identified matches and mismatches between the actual coverage of application scenarios in the literature and the users’

actual privacy demands. In the course of analyzing the personalization usage models, future privacy-related research opportunities in longitudinal and walk-by personalization were identified. This section also pointed out privacy threats that may be interesting to look at in the future, i.e., repudiation and decontextualization, as there appears to be little coverage so far. The section then presented a heat map of 25 countermeasures that can be applied to each threat category. Eventually, the section discussed the research opportunities identified in the course of the survey: remaining analyses of privacy threats (i.e., tampering and decontextualization) in particular application scenarios (i.e., social grooming, demonstrating achievements, current and past work processes, and planning and information overview).

7.6. Toolkits and Frameworks

According to Davies et al., “pervasive [public] display research is still in its infancy” [62]. The temporal distribution of publications presented as part of the survey conducted by Ardito et al. [15] further underpins this claim: Prior to the year 2000, their classification scheme matches only 14 papers in total; between 2001 and 2014, however, even individual years trump this number. This infancy is probably most noticeable with regard to established research methodologies and tools:

Pervasive displays represent a young and exciting area of research. In many areas of computer science there are well-accepted research tools and techniques that are used to help answer common research questions. [...] For researchers in pervasive displays, however, the choice of tools and techniques is much less obvious—there are no widely accepted test data sets, tools or techniques. —*Davies et al. [62, p. 69]*

A number of researchers acknowledged this lack of tools and techniques. They thus worked on generally applicable toolkits for public displays that can be employed in various research projects. The remainder of this section first focuses on these (tangible) toolkits before addressing (more theoretical) frameworks for public displays.

The *iScreen toolkit* introduced by Handte et al. [91] is designed to support researchers in developing display prototypes that involve touch-centric applications. They found that researchers often had to customize integrated computer systems and their software to accommodate the needs of a specific research scenario. This includes the design of suitable user interfaces, “glueing together” various pieces of software, e.g., databases and front-ends, as well as protective software, that restricts access to functions provided by the operating systems, for example. Consequently, their *iScreen* toolkit “[...] aims at minimizing the development effort by providing a set of reusable building blocks for interactive applications” [91] in a generic and extendable way. The toolkit comprises various user interface components, e.g., a touchable keyboard, an HTML browser, or a calendar viewer, and means of interaction, e.g., a camera or Bluetooth connectors.

A similar approach are *PuReWidgets* as proposed by Cardoso and José [45]. This toolkit comprises a set of widgets, i.e., user interface elements, and services that handle the user input via multiple means of interaction. The authors aim at “[...] the creation of a programming toolkit that developers can incorporate into their public display applications [...]” [45] just as they would when developing a web application based on external frameworks or libraries. This is achieved, for example, by providing a high-level abstraction layer that hides the technical characteristics of a specific interaction technique, e.g., Bluetooth, from developers or researchers.

The *Proximity toolkit* by Marquardt et al. [133] is built around the principles of proxemic relationships between users and public displays. Examples of such relationships would be the user's distance and orientation towards a display. Marquardt et al. note that though proxemic interaction appears to be a very common way of human behavior, "only few ubiquitous computing (ubicomp) systems interpret such proxemic relationships" [133]. They found that one reason is rooted in the technical difficulties developers and researchers experience while developing public displays that support proxemic interaction: A plethora of sensors, interfaces, and APIs render the development a challenging task. The Proximity toolkit supports this process with its components, for example, "tracking plug-in modules," a "visual monitoring tool," or an event-driven application programming interface. This way, the toolkit allows for rapid prototyping of public displays that incorporate proxemic relationships.

The *SenScreen* toolkit by Schneegass and Alt [200] focuses on sensor-enabled multi-display networks. Comparable to the PuReWidgets and the Proximity toolkit, SenScreen strives to encapsulate the technical details of various sensor sources, e.g., Microsoft Kinect, in a high-level API. The aggregated and enriched sensor data is distributed to multiple clients, e.g., public displays, via a central server. This allows multi-display networks to react appropriately to user input, for example, in cases where "multiple applications access the same sensor on a display or where games run across multiple displays" [200].

While the previously presented toolkits provide general means to developers and researchers of public display, the approach presented by Dang and André [61], called *Environs*, aims at very specific application scenarios: "multi-display environment applications supporting interactive real-time portals" [61]. Real-time portals show (parts of) a public display and let users (remotely) interact with the content shown on

the portal or the public display, respectively. Smartphones or tablets may be manifestations of such real-time portals. As there are many devices, resulting in various hardware as well as software specifications, the authors identified the need for a homogenous development environment. Their approach thus “helps developers and designers focus on application and presentation logic” [61] rather than on individual engineering issues.

The term framework is very often used in literature about public displays. Many publications refer to a framework when describing a general (theoretical) structure. This section, however, focuses on those publications that propose a concept, which is actually applicable to other public display systems. The following subsections break down the related work into four parts: (i) design, (ii) content, (iii) interaction, and (iv) social connections. Each subsection presents corresponding frameworks in more depth.

7.6.1. Design

The paper by Vogel and Balakrishnan about *interactive public ambient displays* [231] is often cited in literature. This is probably due to the fact that it is an early publication in the field of public displays that covers fundamental design aspects and introduces novel interaction concepts. One example is the interaction framework, that is based on four phases: (i) ambient display, (ii) implicit interaction, (iii) subtle interaction, and (iv) personal interaction. This interaction framework addresses “[...] sharable, interactive public ambient displays that support the transition from implicit to explicit interaction with both public and personal information” [231]. In a prototypical implementation, they present techniques for each interaction phase so that designers

and researchers may employ these (partially futuristic) proposals for user interfaces and means of interaction in their own projects.

Memarovic et al. propose “a layered framework addressing the multifaceted issues facing community-supporting public display deployments” [145]. As the title implies, this framework focuses on public displays and is thus called *P-LAYERS*, which stands for “public display layers.” As most of the framework authors tend to evaluate their research on public displays in deployment-based studies [7, 9, 50, 52, 143, 224], they regard themselves as versatile and experienced in that context. To help other researchers and designers address the challenges that arise when deploying public displays “in the wild,” they developed the five-layered framework as depicted in Figure 7.4.

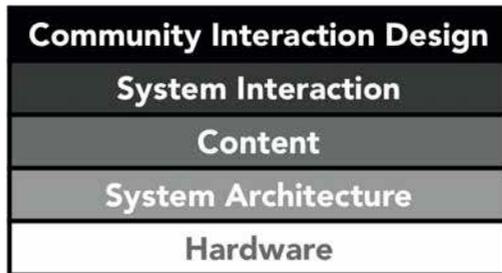


Figure 7.4.: The P-LAYERS framework as presented in [145].

The vertical arrangement implies a sequential process that should be completed from bottom to top: (i) The hardware should “[...] fulfill requirements and expectations, both from users and the researchers” [145]; (ii) the system architecture should consider the interactivity and durability of the deployment; (iii) the content should be suitable—also technically—and stem from appropriate sources; (iv) the system interaction should cover three questions about “where to place the display,” “which level of complexity is appropriate,” and “how [...] interac-

tion [should] be triggered” [145]; (v) the community interaction design should ensure that users “[...] understand the meaning of the application” [145]. The framework may thus “provide useful support for researchers/practitioners by alerting them to the diverse range of issues they are likely to encounter with the development, deployment, and maintenance of public display systems that aim at simulating community interaction” [145].

The *APEX* framework presented by Silva et al. [206] is an approach towards model-based rapid prototyping of ubiquitous environments. It enables designers, researchers, and users to experience an envisioned system in a 3D simulation. The framework is based on three components: (i) the virtual environment, (ii) the behavioral component, and (iii) the communication/execution. The underlying model of the framework is based on a Colored Petri Net (CPN). This way, the framework allows for a very precise simulation of the ubiquitous environment. Yet, the creation of such a virtual environment may be a complex task, e.g., modeling all involved devices, and may thus demand too much of “average” designers or researchers. An opposite approach is based on miniature models as proposed by Nakanishi [159]: Miniature models of real locations are used to facilitate rapid prototyping. However, not all relevant aspects of an ubiquitous environment can be addressed in a miniature model, which is why the author suggests to analyze a corresponding virtual model as well. This virtual model can then be used, for example, to assess the ideal positioning of interactive devices, while the miniature model can be used to eliminate discrepancies between the virtual and the real space, e.g., optical attenuation. The approach may be subject to some limitations, as the virtual model and the miniature model may lack realism. Additionally, in terms of handiness, the system requires designers and researchers to work with two models rather than just one.

Calderon et al. [42] introduce *RED*—the Really Easy Displays framework. Their approach employs web technologies, e.g., JavaScript, to rapidly prototype multi-display applications. Though RED focuses on the interoperability between multiple displays with regard to content and interaction, it “can be limiting for applications outside the common scenario of a mobile device interacting with situated displays” [42]. Similar to some toolkits presented above (cf. PuReWidgets [45], Proximity toolkit [133], or SenScreen [200]), RED aims at higher-level abstractions of various data streams and interaction modalities.

Ardito et al. propose a three-part framework for “highlighting factors of interest and organizing a systematic approach to the evaluation” [16] of public displays: (i) environmental factors, “which consider the physical location where the display is installed (e.g. city street, museum hall, fair, train station). These factors include how the display is positioned in the environment, for example whether it is in a very visible location, in a crowded place, etc.” [16]; (ii) software factors, “such as interface design, functionality, speed of processing” [16]; (iii) hardware factors, “such as the technology adopted, screen size, screen orientation (tabletop vs. wall screen)” [16].

7.6.2. Content

Rogers et al. [187] present *BluScreen*, an auction-driven approach to content delivery on public displays. The framework focuses on advertisements, that can be targeted at individual users. The system is able to identify users based on the presence of their Bluetooth devices, e.g., smartphones. A comparison of different “bidding strategies” indicated, that their approach surpasses other (basic) approaches and may successfully “predict the arrival and departure of users” [187]. Though

BluScreen is specialized on advertisements, it could also be used to efficiently schedule other types of content.

PEACH is also concerned with individual content in a very specific application scenario. It is a “framework for museum visits, focusing on aspects where adaptivity is central” [216]. *PEACH* addresses mobile devices as well as public displays with a set of interactive and adaptive components, i.e., (i) animated agents, (ii) adaptively generated videos, and (iii) generated post-visit reflections. The authors focus on mobile devices, appropriate user interfaces, and ways of influencing the autonomous decisions of the system for future museum visits.

Another approach to realize context-aware adaptation of display content is presented by Cardoso and José [44]. As mentioned in Section 7.1, their framework uses digital footprints “to deliver ‘the right information at the right time’” [44]. Footprints may comprise all traces that users create while interacting with public displays, for example, by pressing buttons or by simply being present at a certain location.

As public displays usually have a large screen—so that it is visible from afar and may be used by multiple users at the same time—arranging and exploiting the available screen real estate is an important task. The *web-based framework for spatiotemporal screen real estate management* as presented by Lindén et al. [130] addresses this task. It allows for “[...] dynamic partitioning of the screen real estate into virtual screens assigned for multiple concurrent web applications” [130]. The framework is used, for example, in the *UBI-hotspots* [164, 165, 166]. A special characteristic of the framework is that it can only accommodate web-based public display applications.

The framework presented by O’Neill et al. [168], see Section 7.4, focuses on showing private data, e.g., patient records, in public spaces—hospital waiting areas in this case. They propose to treat users as citi-

zens, as system designers are not able to anticipate who exactly would use their system:

We may know little or nothing about the users of a publicly available, large-scale pervasive system, but there are a number of things we can know about citizens. Such information includes citizenship rights, how citizens view public systems (e.g. TV, public transport etc), and what type of access to public systems citizens prefer or require.

—*O’Neill et al. [168]*

This notion of citizens demarcates the first dimension of their framework. The second dimension comprises three information spheres, i.e., (i) public, (ii) social, and (iii) private spheres, cf. Section 7.4. The third dimension of the framework consists of spaces, which are in line with the dimension of information spheres: (i) “public spaces are open to everyone, mainly because they usually belong to the community itself, e.g. a town square is a public space” [168]; (ii) “private spaces are spaces controlled by an individual, which can be used in whatever way the owner sees fit. Private spaces promote a sense of security and privacy, such as a bedroom or a toilet” [168]; (iii) “social spaces are those spaces that are neither private nor public. Examples of such spaces include homes, cars, hospital treatment cubicles etc.” [168]. According to O’Neill et al., this “framework will not design systems for us but it will support us in making design decisions” [168].

Finally, the *WE-BAT* framework presented by Elhart et al. [74] addresses the issue of scheduling interactive and conventional applications, such as advertisements, on public displays in a reasonable way. Sometimes, public display stakeholders require content to be shown in a particular order or within a specific (temporal) context. Thus, Elhart et al. identified common scheduling constraints and other requirements that

contribute to their framework. The framework also provides a formal notation and an API that developers may employ to build future modular web-based public display systems.

7.6.3. Interaction

The audience funnel as introduced by Michelis and Müller [146] is a framework often referred to in literature on public displays. It contains elements related to the design of public displays in general, see Section 7.4, and elements that address the interaction between users and public displays. The framework differentiates between the six interaction phases as follows: (i) “everyone who happens to be present in a certain vicinity of a public display can be called a *passer-by*” [146]; (ii) “as soon as a passer-by shows any observable reaction to the displays, such as looking at it, smiling or turning his head, he can be considered a *viewer*.” [146]; (iii) “as soon as the viewer shows any signs of movement that is intended to cause some reaction by the display, we can call him a *subtle user*” [146]; (iv) “[...] after some initial subtle interactions users usually tried to position themselves in the center of the display [...]. Such a user can be called a *direct user*” [146]; (v) “many users started to interact with the other displays after a phase of direct interaction with one display. Such a user can be called a *multiple user*” [146]; (vi) “[...] many users conducted *follow-up actions* [emphasis added] after direct or multiple interaction. For example they took photos of themselves or their friends while interacting with the displays [...]” [146].

The *peddler framework* proposed by Wang et al. [235] builds upon the audience funnel by Michelis and Müller [146]. Wang et al. incorporate three additional aspects: (i) continuous, fine-grained proxemic measures (e.g., the user’s distance and orientation), (ii) reacquiring the

user's interest after phases of waning attention, and (iii) attentional states, that relate to the user's short-term interaction history. The objective of the peddler framework is "to lead the passerby into a more attentive stage" [235]. This objective is aligned with the *AIDA* strategy, i.e., to attract *attention*, maintain *interest*, create *desire*, and lead customers to *action* [235]. Thus, the most apparent difference to other approaches is, that the peddler framework "models the *display's* goals as opposed to the *user's* goals" [235].

The crossmodal display framework, as used in CrossFlow and CrossBoard, by Cao et al. was already described in Section 7.2. However, with regard to interaction, this framework allows multiple users to interact in parallel since it draws on the users' smartphones for individual information delivery. Furthermore, the authors emphasize the importance to balance the visibility of a public display and the ease of accessing the shown information. Thus, CrossFlow and CrossBoard make use of crossmodal cues to point out relevant information to individual users of public displays in an easily perceptible manner.

Another approach, also based on personal mobile devices, i.e., smartphones, is *BlueTone* as proposed by Dearman and Truong [69]. *BlueTone* uses audio transmissions via Bluetooth to control public displays. The authors identified their system to be vulnerable to spoofing, as attackers could manipulate the hardware components and pretend to be the public display. Users would then send (personal) data to the attackers rather than to the actual public display. Nevertheless, *BlueTone* is designed as "[...] a framework that supports opportunistic interaction with public displays using any Bluetooth enabled mobile phone" [69]. Dearman and Truong focus on this type of interaction, as it allegedly allows for remote multi-user interaction without the need of dedicated software on the users' smartphones.

The *ContentCascade* is an interaction method based on the content exchange framework presented by Raj et al. [184]. It allows to transfer short summaries of contents from public displays to the users' smartphones. This way, the system may help users to remember the seen information when reading the summaries later on. Users trigger the information transfer either explicitly, e.g., by clicking buttons, or implicitly, by hovering in front of the public display, for example.

The framework introduced by Baldauf et al. focuses on “the camera-based control of large markerless displays through smartphones in real-time” [19]. Similar to the crossmodal display framework, Blue-Tone, or the ContentCascade, this interaction method is also based on smartphones. The framework incorporates fundamental aspects of *augmented reality* (AR), as it uses them to overlay virtual buttons on top of live video footage of public displays. Users may thus use their smartphones as a *Magic Lens* [28] to augment their visual perception of such displays. For example, this framework allows users to press buttons on their smartphones rather than on the public display itself to trigger actions. This way, the framework allows for remote interaction of multiple users in parallel. Similar to SenScreen [200], Environs [61], and RED [42], Baldauf et al. designed their framework to be applicable in multi-display networks.

This subsection presented an overview of related work with regard to interaction frameworks. Apparently, researchers have an inclination to use mobile devices, e.g., smartphones, to facilitate interaction between users and public displays. As Dix and Sas note, “it is increasingly rare to find someone without a mobile phone [...]” [73]. Dix and Sas thus looked at possible synergies and opportunities of this frequent combination of smartphones and public displays. They propose a framework that can be “used to analyse potential issues, problems and requirements in particular those involving conflicts between

individual interaction and ‘audience’ experience, where the audience includes passers-by and bystanders around the public screen” [73]. Their framework analyzes existing design spaces and also contributes their own design space, which partially draws on concepts introduced above, e.g., the notion of spaces by O’Neill et al. [168]. Furthermore, their framework presents two design strategies: (a) to hide personal interaction on the one hand side and (b) to expose interaction on the other hand side.

The first one is especially interesting in terms of privacy as “one of the most obvious uses of the personal device is simply to use it to display parts of the content or interaction that we do not wish others to see for reasons of privacy or intrusion” [73], cf. Table 7.9. Additionally, this design strategy can be used to (a1) avoid information overload, (a2) reduce resource conflicts, (a3) restrict content, and to (a4) weave “personal choice into public schedule” [73]. The second design strategy is to “deliberately expose the effects of individual interaction” [73]. This strategy may help to (b1) negotiate control, (b2) select the audience, (b3) ensure accountability and auditability, as well as to (b4) create enticement and engagement, which is similar to (b5) creating an experience or performance that can be perceived (and possibly enjoyed) by others.

7.6.4. Social Connections

As their name implies, public displays are usually installed in places that are accessible to a broader audience, e.g., shopping malls, civic centers, or airports. Naturally, people are likely to interact with each other in such places—either explicitly or implicitly. Some researchers thus analyzed aspects of public displays with regard to, for example, social awareness, sparking conversation, or increase social cohesion.

In doing so, many publications focus on one specific research objective or one particular implementation rather than on a generic framework for social connections:

PlasmaPlaces [56], and in particular *CHIplace* or *CSCWplace*, as presented by Churchill et al., is a large web-based public display, designed to raise people's awareness of activities in online communities in real, physical places. Their system is designed to spark and foster interaction between participants of scientific conferences. They identified two main principles that support this objective: (i) "cooperation is the main behavioral intent that must be supported" [56]; (ii) "identity and social interaction are key elements underlying cooperation" [56]. According to their work, "these two elements can be supported via social, technical, and socio-technical means" [56]. *PlasmaPlaces* allow users to quickly navigate within the available content that has been collected from various sources, e.g., websites or blogs. After browsing the content, the system also encourages users to perform follow up actions (cf. the audience funnel in Figure 7.1) by pointing at further community resources. Churchill et al. also present general findings with regard to the social surroundings of public displays: (i) "People respond positively to faces and other indications of community member identity, including names and contact information" [56]; (ii) "people are attracted to large central displays as a focus of attention" [56]. Similarly, the *Proactive Displays* introduced by McDonald et al. [141] or the *IntelliBadge* presented by Cox et al. [60] facilitate social exchange among conference attendees.

Besides these application scenarios, which are rooted in a scientific context, there is also research about the social characteristics of public displays in other domains. *MobiComics* [131], for example, is designed to stimulate social interaction between people, who collaboratively work on a comic strip. A field study showed that Mobi-

Comics actually managed to spark interaction between users, while it also raised concerns about the users' privacy at the same time. *My-Position* [230] is supposed to initiate civic discourse opportunistically. Just like MobiComics, the system was also evaluated in a field study: Apparently, users actually prefer that others can see how they vote, since this may help to engage a debate. Yet, this result likely depends on the individual application scenario: People might be more conservative when casting a ballot on presidential elections, for example. Finally, Böhmer and Müller propose *social signs* [32], an approach that envisions public displays to autonomously draw on information from social networks and to visualize this information. This way, the public displays would be able to highlight unexpected or opportunistic links between passersby, which could, in turn, increase social cohesion. According to an online survey, most users would prefer to publicly show their real names, their interests, their contact options, and friends they have in common with other users. Yet, their study also points out that some participants raised serious concerns about their privacy: "Interestingly, the participants' free text answers do not correspond with our quantitative data. Most of them negatively criticize the social signs due to privacy and security" [32]—this supports Nissenbaum's observation of an apparent paradox of people's demand for privacy, cf. p. 76.

Besides these publications, that are concerned with specific public display installations, there is also research on more generic frameworks with regard to social connections. The remainder of this subsection presents two examples. Brignull and Rogers [38], for instance, propose a framework that provides design implications for interactive public displays. Their work draws from the opportunistic observation that most passersby apparently hesitate to interact with public displays, due to, e.g., social embarrassment. Their approach strives to mitigate possible (social) barriers and to encourage passersby to inter-

act. The framework by Brignull and Rogers has been implemented in a prototype called *Opinionizer* [38]. As mentioned in Section 7.4, Brignull and Rogers introduce three activity spaces around public displays. In their conceptual framework, they point out that it is vital to support people in transitioning between these activity spaces: “In particular, in crossing the threshold from peripheral to focal awareness activities (e.g. from chatting to someone on the other side of the room to deciding to move within view of the display to have a better look), people need to be motivated” [38]. Eventually, the framework suggests to present these five key characteristics of an interactive public display to potential users in a direct and intuitive way: (i) “how long an interaction takes” [38]; (ii) “what they will get out of it” [38]; (iii) “what steps are involved” [38]; (iv) “if it will be a comfortable experience” [38]; and (v) “if there is a quick let out, where they can walk away gracefully, without it disturbing the ongoing public activity” [38].

Memarovic et al. introduced the conceptual framework of *Interacting Places* [144]. The framework focuses on “challenges and possibilities for networked public display applications that aim at” [144] stimulating community interaction and place awareness (CIPA). Therefore, it concentrates on four design aspects: Firstly, (i) content providers are comparable to the stakeholders proposed by Alt et al. [5, 7, 8], see Section 2.2.1. Besides content generated by people, for example, social network posts, the definition by Memarovic et al. explicitly includes content generated by services, such as weather forecast websites. Secondly, the aspect of (ii) *content viewers* includes the three categories of unknown group of people, known group of people, and individuals. Thirdly, the distribution of any kind of information is, according to Memarovic et al., always delivered through (iii) *communication channels*. Their notion of such a communication channel is driven by a technical understanding as they refer to, for instance, e-mails or instant messaging. The framework further differentiates between

inclusive and exclusive channels. The first type of channel is “open-for-everyone” [144], while the latter type may be used, for example, to deliver directed messages to specific recipients. In both ways, content can be either targeted at people, e.g., individual users, or places, for example, an entrance hall. Finally, an (iv) *awareness diffusion layer* “describes how community awareness building happens both *explicitly*, i.e., through content tailored towards a specific audience, and *implicitly*, by observing output for other people” [144]. To further underpin this last (possibly abstract) design aspect, Memarovic et al. present the following example:

[...] While foreigners might not be able to understand that “Barca” refers to a football club, or even a sports club altogether, they might still realize that its community is very active in a place due to the number of messages posted bearing the “Barca” logo. Similar implications may be drawn from the artwork and typography associated with the communication: a visitor to a bar may not understand who is posting what on a screen, but might perceive the design as either very professional or very homely, thus getting a sense for a very professional or very caring community, respectively. —*Memarovic [144]*

Based on these four design aspects, the conceptual framework presented by Memarovic et al. may be used to analyze existing public display systems with regard to social connections, as well as to support the design of future installations.

8

Summary

This part laid out the scientific foundation of the thesis. Chapter 5 explained the design and use of the applied descriptive as well as experimental research methodologies, see Table 5.1. Key concepts in public display systems were presented in Chapter 6. There, the first section emphasized the importance of a common understanding of the term context. Accordingly, the definition used throughout this thesis was presented and located within related work. The next section examined existing notions of privacy. “One point on which there seems to be near-unanimous agreement is that privacy is a messy and complex subject” [161, p. 67]. Yet, it is important to establish a common ground and understanding for further analyses and discussions, as presented in the course of this thesis. Furthermore, a naive approach to privacy based on a public/private dichotomy may be unsuitable in most public display application scenarios. To present more nuanced approaches to privacy, the section introduced Solove’s taxonomy of privacy [212] and Nissenbaum’s framework of contextual integrity [161]. Finally, the section tried to debunk the often used “I’ve got nothing to hide” argument by showing that privacy is a substantial component of people’s everyday life, because it allows them to act and evolve freely while their individual personality remains untouched.

The following section explained the term personalization by locating it within related work, highlighting different design perspectives, and presenting possible sources of personalized data. The subsequent section presented design aspects of public displays while considering hardware as well as software components. The section also looked at established (design) models in general and pointed out that there are no such commonly accepted models with regard to public displays.

The next section introduced the concepts of threats, threat models, and countermeasures. In the course of this section, the seven threat categories addressed by STRIDED* were presented and explained. STRIDED* is the privacy threat model that constitutes the first scientific contribution (C1) of this thesis.

Chapter 7 located the key concepts defined in the previous chapter within related work. Section 7.1 presented examples of context-aware public display systems. While presenting existing publications on privacy, Section 7.2 highlighted the fact that though privacy is regarded as important and valuable in theory, the general public appears to feel indifferent about it at best. This ostensible paradox or contradiction should be considered when designing privacy-preserving personalized public display systems. The next section highlighted different aspects of personalizing public displays and presented the personalization usage model as proposed by Davies et al. [64]. The following section discussed publications concerned with the design of public displays. At the beginning, it referred to publications identifying and addressing particular challenges that arise when designing public display systems. This is followed by a comprehensive overview of existing design considerations, design spaces, and design analogies, such as the audience funnel or the honeypot effect, for example. The section concluded with the “taxonomy of visual display-based activity in office spaces” as presented by Perry and O’Hara [183].

Section 7.5 presented another major contribution of this thesis: an extensive literature survey. The survey comprised 120 publications between the years 2000 and 2014. The results provided new insights into two domains: (i) the distribution of application scenarios and the types of personalization being applied; and (ii) the distribution of privacy threats as well as proposed countermeasures. The survey was based on a novel classification scheme based on established concepts; the scheme proved to be valid and useful to derive the findings. With regard to public displays, the results highlighted (i) most prominent application scenarios (e.g., quick reference), (ii) most common personalization usage models (e.g., active personalization), (iii) most imminent privacy threats (e.g., information disclosure), and (iv) most often used countermeasures (e.g., to use a 2nd device). Additionally, the list of existing countermeasures was also presented as a heat map. For each STRIDED* privacy threat, this heat map allows designers and researchers to quickly find the countermeasure, that has been used most frequently—and that may thus be most suitable.

Chapter 7 concluded with an outlook of existing toolkits and frameworks. This outlook was subdivided into four parts, each focussing on (i) the design of public displays, (ii) the composition of shown contents, (iii) the usage of interaction modalities, and (iv) the potential of social connections. Section 7.6 points out that “pervasive [public] display research is still in its infancy” [62] and that “there are no widely accepted test data sets, tools or techniques” [62] for researchers focusing on public displays. This further emphasizes the significance of the third scientific contribution (C3) of this thesis, i.e., a process integration that comprises a methodology and tools to support the design, prototyping, and evaluation of privacy-preserving personalized public display systems: C3 and the other presented toolkits or frameworks may catalyze future public display research and designs.

III

Designing Privacy-Preserving Personalized Public Display Systems

9

Challenges

As explained in Section 7.4, research on public displays is a young discipline. This unfortunately implies that researchers may not resort to established methods, models, or tools. Davies et al. [62] suppose that the lack of a model for public displays is one reason for why research on public displays remains challenging. The results of the extensive literature review presented in Section 7.5 allowed to identify common design challenges that may serve a first approach towards such a model for public displays. The related work was clustered manually according to the focus of each publication: Aspects that were portrayed as challenges, issues, or central points were extracted, consolidated, and finally used to derive the eight challenges described in Sections 9.1–9.8. As public display systems consist of at least two components, i.e., software and hardware, it is vital that a holistic model accounts for both components, cf. Section 6.4.

Naturally, this list is not final, as further challenges may exist. However, it likely includes major aspects with regard to the design of public display systems. Future work may extend this first approach towards a more comprehensive model for public displays. Section 7.6 presented existing toolkits and frameworks that support the design process of public displays systems. However, some of the challenges identified

by the model presented below are only partially addressed. This thesis thus proposes a particular process integration (C3), see Section 10.3, as an attempt to bridge this gap.

Finally, the challenges presented in the remainder of this chapter may contribute to an understanding of the context of a public display system, cf. Section 6.1. Each challenge may be regarded as a dimension of a design space that guides the process of identifying “any information that can be used to characterise the situation of an entity” [72]: This open and inclusive definition of context by Dey may be ambiguous or of no avail in some situations. Designers of public display systems, for example, could thus use the challenges as a guide. With regard to Nissenbaum’s (social) notion of context as introduced in Section 6.1, the challenges may also contribute to her framework of contextual integrity: Dynamic environmental factors (see Section 9.4), user acceptance (see Section 9.7), and legal constraints (see Section 9.8) may have an impact on her context-relative informational norms.

9.1. Situatedness

Public displays are always situated in a certain location [65]. Ojala et al. emphasize that it is key to pinpoint ideal locations: “We have discovered that location is central to the way people use [...]” [164] public displays. Likewise, Davies et al. note that “in addition to tailoring content to specific users, considerable value can be obtained from ensuring displays evolve to provide content appropriate to their situation. Ensuring the appropriateness of displays [...] is a key challenge” [62]. Other publications second this opinion, e.g., Huang et al. [98], Snowdon and Grasso [211], Alt et al. [8], or Perry and O’Hara [183]. Clearly, the situatedness of public displays defines their physical location, but in most cases it also impacts the shown contents and

offered services. For example, public displays installed at train stations usually show content related to rail traffic, such as arrival and departure times, or their contents are designed to entertain waiting passengers. This characteristic constitutes a significant advantage of public displays over location-based services on personal devices such as smartphones, for example. While such versatile devices may be used in a plethora of scenarios by installing according apps, they may also have some downsides.

Firstly, users have to be made aware of the (per se invisible) location-based services, especially in case of first time users, as their smartphones may not be configured yet. Moreover, users may refrain from installing and using dedicated apps or applying special configurations for various reasons, such as technical inexperience or potential costs, like roaming fees, for example. Secondly, information provided by a smartphone is not available instantaneously: Users have to take out their smartphone, turn it on, enter a code to unlock it, find and start the corresponding application, wait for the application to load, and finally navigate within the application to find to the desired information. Thirdly, the use of smartphones may be inappropriate or impossible in some situations, for example, when performing religious acts [90] or due to insufficient network coverage in subway stations or remote rural areas [223, 224].

Often, prototypes of public displays have to be deployed at the intended locations to truly analyze the situatedness of the final system. Such field studies have certain advantages, for example, a high ecological validity [12, 71], but also some drawbacks [71]: (i) The context and conditions are not as controllable as in a lab, especially in terms of repeatability; (ii) deployments require robust prototypical implementations that may operate independently for a certain period of time [217]; finally, (iii) there may be organizational overheads, such

as the need to transport the apparatus and material—including developers, users, or study participants—to the deployment site. Instead of field studies, designers and researchers may resort to lab-based evaluations. Yet, this approach also has some disadvantages to it, for example, the lack of realism, especially with regard to environmental or societal aspects [71]. As suggested by Delikostidis et al. [71], the individual optimum may lie somewhere in-between both extremes. The challenge is thus to strike a balance.

With regard to the physical location of such displays, designers may ponder different options to optimize the flow of huge crowds, especially in case of an emergency. The *EyeCanvas* [57] project compares public displays in semi-public locations with installations in an open community space. It also discusses the process and importance of finding an appropriate installation site. The *GAUDI* system is an indoor navigation system based on small public displays that are aware of their individual location and may react accordingly if relocated. [117, 119]. The PEACH project [216] provides historic information to museum visitors according to their location within the exhibition. *SPAM* [50] carefully integrates the location of the display in the design process, as the system may display sensitive patient data. The situatedness of a public display is also addressed by the APEX framework [206], the miniature models by Nakanishi [159], the concept of digital footprints by Cardoso and José [44], the information spheres by O’Neill et al. [168], PlasmaPlaces [56], and MyPosition [230].

9.2. Form Factors

As diverse as the application scenarios for public displays are the form factors that go along with that diversity: Their sizes range from small watches—actually used as public displays [179]—to huge facades [82].

The latter ones may be so large, that users may not perceive the entire display at once. The results of the field study performed by Huang et al. [98] indicates that “[...] people seemed to linger at smaller displays for a longer period of time” [98]. Their results also provide some insights with regard to interaction and privacy:

The use of a smaller display may also create a more private or intimate setting within the greater public setting that leads a viewer to feel less exposed and therefore encourages a longer interaction and greater comfort with displays within a public space. —*Huang et al. [98]*

Consequently, the form factor should match the particular application scenario to optimize the experience for the user. A public display, for example, that should be used by multiple users in parallel to browse photos should probably be larger than a regular desktop monitor. Small displays may, however, provide some privacy when handling sensitive personal data. Hence, the challenge is to carefully define the form factor as early as possible, in particular as this physical characteristic may not be changed easily in subsequent research or design phases.

The *CityWall* [180] analyzes the influence of the size of the system on the audience’s interactional behavior. The design of the *WrayDisplay* [222] incorporated the opinion of actual users about their preferred form factors. From a slightly different perspective, SPRIOS [185] as well as the approach by Brudy et al. [40] use varying (virtual) screen sizes to enhance the user’s privacy and to protect sensitive personal data from the glances of passersby. With regard to frameworks, the APEX framework [206] addresses this challenge, as well as the miniature models by Nakanishi [159] and the evaluation framework proposed by Ardito et al. [16].

9.3. Fixed Environmental Factors

Storz et al. [217] present a number of lessons learned from three long-term deployments of public display systems. The summary of one lesson is that “environmental challenges can be significant” [217] and that one should “never underestimate the impact of environmental factors on a deployment” [217]. These factors include, for example, objects such as furniture, vegetation, or buildings. Fixed environmental factors are thus closely related to the first challenge addressing the situatedness of public displays. The emphasis, however, is different: While the first challenge focuses on the purpose of the display in terms of content and services, this challenge points to the relation to and possible issues with existing physical objects, cf. the evaluation framework by Ardito et al. [16] as presented in Subsection 7.6.1. For example, with regard to large-scale public displays, such as media facades, there may be a sweet spot [176] due to technical or architectural constraints. It is paramount to locate this sweet spot as early as possible in the design process, so that it will not interfere with other fixed environmental factors, e.g., buildings or streets. Besides media facades, Huang et al. [98] suggest that this challenge also applies to other public displays:

[...] The vast majority of large displays in public areas were designed [...] with less an a focus [sic] on how people would be moving within a space and how other activities within or aspects of the space might affect use of the display. [...] The ultimate position and context of the display should be taken into account during the design phase. —*Huang et al. [98]*

With regard to interaction, the analysis of fixed environmental factors is also recommended to guarantee the accessibility of means of interaction, e.g., keyboards or touchscreens, as they might be blocked

by other objects. Accordingly, Memarovic et al. claim that “the location and exact placement significantly affects how users approach and interact with a display” [145]. In a similar vein, Huang et al. suggest to consider “how the surrounding environment can be designed or taken advantage of to draw attention to the displays” [98]. It may be challenging, however, to define which fixed environmental factors are relevant and which might be neglected. Furthermore, it may be difficult to incorporate the identified factors appropriately in research or design processes.

During the evaluation of the *MobiDic* public display system, Müller et al. [153] reported on the observation that other objects, such as kicker tables, could prevent passersby from using displays. In contrast to this, Steinberger et al. [215] exploit the fixed environmental factors that their public display voting system is exposed to in order to foster user interaction. Further projects that consider other environmental factors are, for example, MyPosition [230] (columns that people tend to gather around), Proactive Displays [141] (coffee tables that conference attendees grab beverages at), or the Opinionizer [38] (a bar ensuring a constant flow of people looking for beverages). Some frameworks also address this challenge, for example, the APEX framework [206], Nakanishi’s miniature models [159], or PEACH [216].

9.4. Dynamic Environmental Factors

External factors such as passersby, time of day, or the weather [164] may induce effects on public displays. Which factors to account for in particular may depend on the specific application scenario at hand.

Weather also affects [...] use, even indoor[s] [...]. After mapping our logs of average daily temperatures and weather conditions (sunny, cloudy, raining, snowing), we found that sunnier and warmer days correlate with higher [...] use in terms of screen touches, services launched, and user interaction time. Our correlation analysis attributes about 10 percent of use variation to changes in ambient temperature alone, discarding other variables such as time of day, day of the week, or even location. —*Ojala et al. [164]*

The essence of this challenge may already be apparent: It may be demanding to identify all relevant factors and integrate them reasonably in research or design processes. Due to the ever increasing plethora of sensing devices and means of interaction, it is impossible to present a comprehensive list of available technologies. This section may rather provide some examples to indicate the breadth of this challenge.

Many public displays react to the user's position or orientation in front of the system. Some design models presented in Section 7.4 also integrate this concept, cf. the subtle interaction phase of the audience funnel by Michelis and Müller [146]. A very common approach taken by many researchers is to turn the user's smartphone into a versatile tracking device. In most cases, technologies such as Bluetooth, WiFi, or *NFC* are used to sense the users, see [6, 108, 138, 140, 203], for example. Frequently, the raw data about the user's presence is enriched

by additional logic in order to make sense at a higher level: The approach based on *adaptive user profiles* as presented by Alt et al. [6], for example, tracks the user's position and aggregates this information with additional data, such as point of sale visits, public display views, event ratings, and coupon redemptions to draw further conclusions.

Similarly to this project by Alt et al., there are other systems that sense and respond to individual users. The Proxemic Peddler by Wang et al. [235], for example, is built around the notion of proxemic interaction and reacts to implicit as well as explicit user behavior. In contrast to this, the approach proposed by Kurdyukova et al. [122] deliberately ignores individual users, but rather analyzes and reacts to groups of people. The authors underpin their design decision with the plausible assumption that most public displays are installed in prominent, well-frequented places. It is hence less likely that individual users will be in the vicinity of the public display, but rather aggregations of people.

Some projects use the sensor data to infer information about approaching, viewing, and leaving users in order to implement privacy. SPIROS [185] and the approach entitled *Is Anyone Looking?* [40], for example, adapt the content shown to individual users as soon as additional users appear. This way, these systems lower the chances for third parties to cast glances at the possibly sensitive data of individual users.

Besides the frameworks mentioned above, there are also other frameworks that account for dynamic environmental factors: PEACH [216] adapts video presentations shown to visitors of an exhibition to their individual preferences. Additionally, the system generates personal "post-visit summaries that reflect the individual interests of visitors as determined by their behavior and choices during their visit" [216]. The web-based framework for spatiotemporal screen real estate management presented by Lindén et al. [130] dynamically adjusts the arrangement of screen content according to the number of active users.

9.5. Mobile Devices

A large body of research focuses on combining public displays and mobile devices, ranging from custom devices [46] to smartphones [15], for example: “As inherently personal devices, mobile phones provide an interesting complement to the public and shared nature of public displays” [106] as found by José et al. This strand of research decomposes into at least two areas: implicit and explicit interaction. In terms of the former one, mobile devices are frequently used to adapt public displays to the user’s location or behavior. Section 9.4 presents multiple examples of such systems. Besides implicit control, some research projects also employ the user’s mobile device as an implicit information storage, see the *Notification Collage* [87], for example. This system allows users to retrieve information distributed by a public display anywhere at an arbitrary point of time.

Similarly, research on explicit interaction between public displays and mobile devices can also be subdivided into controlling and receiving. Using smartphones, for example, to control public displays may have some advantages over other means of interaction: Firstly, users are familiar with mobile devices. For example, most people use their smartphones frequently on a daily basis, i.e., about 60 minutes per day [31, 73]. Secondly, display systems relying on this means of interaction do not have to expose public interfaces, such as keyboards or touchscreens, that may be subject to *vandalism* or may raise *hygiene* concerns. Thirdly, other means of interaction, e.g., gestures or voice control, may induce privacy issues. It may be inappropriate, for example, to have users enter sensitive data, e.g., passwords, by using salient gestures or their voice—both easily observable. The use of personal mobile devices, as proposed by De Luca and Frauendienst [67], for example, may thus alleviate the threat of information

disclosure. Looking at these advantages, it appears natural that numerous research projects let users explicitly control public displays [20, 33, 131, 218] or media facades [34, 30, 83] via mobile devices.

Finally, users may use these devices to explicitly receive information provided by public displays. The Rotating Compass, CrossFlow, or CrossBoard, for example, issue tactile cues that convey navigational or spatiotemporal information to the user. *Multipleye* [169, 172] is based on the concept of visual multiplexing. Users can use their smartphones as a demultiplexer to select which information to show and to access personalized multimedia content on public displays concurrently. Similarly, *QR codes* can be used to let users receive data by visual means only [58, 176, 172]. With regard to privacy, mobile devices may also be beneficial when receiving sensitive data from public displays. The approach presented by Berger et al. [26], for example, uses wearable devices to let users securely receive sensitive data that has been scrambled in the original content shown on the public display.

Despite these advantages, the interaction and communication between public displays and mobile devices appears particularly challenging due to two reasons: (i) The communication stack that is required to facilitate the actual interaction, e.g., via WiFi or Bluetooth, adds some extra complexity to the software system and the user interface; (ii) the great diversity of mobile devices, in terms of, e.g., various operating systems and screen sizes, can make the testing process a daunting task.

9.6. Multi-Display Networks

Multi-display networks introduce unique challenges to researchers as well as designers [65, 217, 225]. In contrast to single displays, for example, multi-display networks offer an extended set of possibilities: If the system ensures continuity and synchronicity between all attached displays, users may benefit from a continuous experience throughout a building, such as a museum [216] or a shopping mall [65], for instance. Testing these large scale deployments may be time consuming and expensive, as deploying and updating the numerous installations can take a significant amount of time. Moreover, the acquisition as well as the maintenance and management costs for display networks can also be a major roadblock [217].

An early system based on networked public displays in an office environment is *UniCast*, *GroupCast*, and *OutCast* [139]. The system spans multiple application scenarios. For example, users may specify their preferences, e.g., hobbies or other topics of interest, by using personal *UniCast* displays; the *GroupCast* displays then show common interests of co-workers closely located to the display in order to spark conversations and foster social contacts.

In the following years, researchers looked at large-scale deployments of multi-display networks. Storz et al. [217] report on their experiences with three long-term public display systems, known as the *e-Campus* system. They compiled a list of “learned lessons” that may help other researchers or designers in implementing similar systems. In particular, they point out that (i) “deployments are costly” [217], (ii) “environmental challenges can be significant” [217], (iii) maintenance efforts should not be underestimated, and (iv) “content is king” [217] at the same time as (v) “content is expensive” [217].

The UBI-hotspots constitute another renowned long-term public display deployment [165, 166]. Ojala et al. [164] also mention the fact that such multi-display networks occasion costs, e.g., for maintenance, which the project needs to cover in some way, for example, by following their approach of showing paid advertisements. Furthermore, they stress the importance of carefully choosing the location of each display within the network, see Section 9.1. In particular, they point to effects induced by the weather—even for indoor installations—and the corresponding social milieu.

Davies et al. [64, 65] present a privacy-aware public display network called Tacita. This system is built around the notion of trust relationships. The idea is that users do not have to trust an uncountable number of display providers, but rather rely on trust relationships established between them and particular *application providers*. The authors argue, that this way, people would be more likely to accept multi-display networks and make use of them. To further support their case, Davies et al. [65] provide two example application scenarios that outline the envisioned benefits of interconnected pervasive displays. However, it remains arguable whether users would actually feel more comfortable with trusting application providers than display providers: In the aftermath of the Snowden revelations in 2013, many users have become more skeptical in general when asked to provide private or sensitive information to computer systems. Moreover, Tacita does not account for compromised network infrastructures.

Some frameworks also address the challenge of multi-display networks: PEACH [216] strives to provide visitors of a museum a continuous and homogeneous experience throughout an exhibition; the fifth phase of the audience funnel [146], i.e., multiple interaction (see Section 7.4) is based on results gathered in field studies:

In a significant number of cases, after exploring the interaction with one screen, users went on to also explore the different effects shown on other screens. For example, when they had passed a first screen without stopping and then started interaction at the second or third screen, they went back to the first screen and started interaction there. Many users repeated this until they had explored all [...] screens.
—*Michelis and Müller [146]*

9.7. Acceptance

Public displays may be either exposed to the general public on the one hand side or dedicated audiences on the other hand side, depending on the particular installation site, see Section 9.1. To increase the acceptance of public displays within a social milieu, it may be beneficial to let the people concerned participate in early design processes. For example, they could determine well-suited locations or contents [75, 164, 222]. Accordingly, Alt et al. [12] found that user acceptance is one of “the most popular questions researchers tried to answer” [12]:

Often used in early stages of the development process, the user acceptance investigates users’ motives and incentives to interact with a display. It can be assessed qualitatively based on subjective feedback, e.g., in focus groups to collect the target group’s view and concerns [11] or quantitatively based on questionnaires [23]. —*Alt et al. [12]*

Based on the MobiDic shopfinder system, Müller et al. [153] analyzed which factors may influence the user’s acceptance of public displays in two field studies. In particular, they looked at these factors: (i) social contexts, (ii) privacy demands, (iii) the visibility of the displays,

and the combined use of (iv) mobile phones. Memarovic et al. conducted a field study to explore the acceptance of *autopoiesic content*, i.e., “self-generative content that is automatically created by matching local context information with regular scheduled information into content that is highly localized” [143]. With regard to privacy, Schmidt et al. state that “user acceptance of technologies that invade privacy is strongly correlated with the perceived value” [199]. Davis [66] identified two factors to have a major impact on people’s acceptance of (software) systems: (i) perceived usefulness or “performance gains” [66], and (ii) “beliefs about performance” [66]. Based on Davis’ work, Huang et al. [100] designed and evaluated the MERBoard. Their evaluation indicates that a third factor might be of relevance: *appropriability*, i.e., people’s competence to gauge whether a public display would be “the right tool” to work on a specific task.

In conclusion, study results emphasize the importance of obtaining the acceptance of all stakeholders, see Subsection 2.2.1, to build successful public display systems [105, 221]. While this may particularly apply to users, other stakeholders should not be missed. Display owners or space owners, for example, may require public displays to show content based on specific scheduling rules [74]. They could thus be more likely to accept systems that cater for this particular need. Similar to the first challenge, see Section 9.1, however, this may be difficult: Assessing the user acceptance likely requires actual (prototypical) systems to be deployed at the actual installation sites. This way, most relevant aspects [71, 164] are likely captured.

In contrast to building and evaluating physical prototypes, Böhmer and Müller [32] based the design of their social signs on the results of a questionnaire. While this may reduce costs, e.g., as no physical system has to be built, deployed, and maintained, the results may exhibit less ecological validity, see Section 9.1 and Alt et al. [12].

With a special focus on privacy, Schaub et al. [196] investigate the user's acceptance towards showing personal calendars on public displays. Their study results indicate, that users tend to accept such systems when given the chance to fine-tune their individual privacy preferences, e.g., by using a smartphone application as in the PriCal system. Schaub et al. also report on an *acclimation effect*, i.e., “participants appreciated the ambient nature of the displays but did not feel compelled to explicitly interact with them” [196].

Similarly, Baldauf et al. [20] assessed the acceptance of their *Video Wall* in a field study. More specifically, they tried to identify the most accepted interaction mode in order to adjust further research in the most purposeful way. When looking at ATMs—a very common and specific type of personalized public displays—De Luca et al. [68] found evidence that the time it takes to authenticate users may impact the overall acceptance of the system.

Especially when designing (long-term) deployments of public displays, it may be advisable to determine the user acceptance before actually deploying the system. Prototypes may be one approach to this. Similarly, Ojala et al. [165] evaluated the user acceptance of the first UBI-hotspots as the foundation of further work. P-LAYERS specifically address the issue of user acceptance: The framework identifies *contextual acceptance* to be key in designing community interaction.

9.8. Legal Constraints

In some situations, public displays may be subject to specific legal constraints, for example, with regard to security, privacy, or contents. Most probably, the context of a public display has an impact on which constrains or regulations may apply, see Sections 6.1 and 7.1. Besides

this dependency on context, the common ramifications of jurisdiction may add to the complexity of this challenge. Sometimes, statutory provisions should be interpreted by experts, e.g., jurists, to avoid any misconceptions or law suits. Mediating between the designers of public displays and those experts, however, may be a challenge itself: While lawyers, for example, may not be able to interpret technical descriptions appropriately, designers may be overwhelmed with juristic terms.

In case of very large public displays, such as media facades, legal constraints may apply more often as these displays are exposed to a broad audience and may thus impact a lot of people [148, 176]. Yet, also smaller systems may be subject to such legal regulations. The SPAM system, as presented Cheverst et al. [50], for example, handles sensitive patient records of people with a history of psychiatric conditions. The authors interviewed the intended users, i.e., staff members, to determine the specific legal constraints that apply in this context. Yet, Cheverst et al. also point out that irresolvable problems may emerge: “We will continue to attempt to meet some of the ethical and moral dilemmas of designing in and for care settings through careful involvement and acknowledgement of users in the design, deployment, use and evaluation process” [50]. In a similar context, O’Neill et al. [168] propose to use public displays in waiting areas at hospitals. Here, the displays may soothe the temper of waiting patients by showing information about expected waiting times and remaining treatments or tests. With regard to medical records, O’Neill et al. identified the classification between private and social information spheres, see Section 7.4, to be challenging: “Indeed, our identification of the relevant information sphere as social may raise ethical or legal questions about the nature of patient records as private or social information” [168].

Also in terms of ethics, Langheinrich et al. [126] present “a step-by-step ethics process that aims at providing structured yet lightweight guidance [...], both stimulating the design of ethical user studies, as well as providing continuous documentation” [126]. Their paper may rather focus on research on public displays than on the deployment of systems ready for the market. Nevertheless, their framework hints at the intricacy and ramifications of the topic.

Furthermore, this challenge may also appear in scenarios apart from harsh or explicit statutory provisions. Hamhoum and Kray [90], for example, looked at how to support pilgrims at religious sites with public displays in unobtrusive ways. Even though there may not be a law restricting the use of public displays in such scenarios, social norms or expected behaviors may be major roadblock.

Finally, Storz et al. [217] also address this challenge in one of their lessons learned from long-term public display deployments. The lesson is entitled “follow the rules” and suggests to “anticipate and plan for regulatory compliance issues” [217].

10

Approaches

The introduction presented in Part I motivated the research questions and scientific contributions of this thesis. The overall objective is to support the provision of relevant content on public displays in order to address display blindness. One promising approach is personalization. However, studying the related work in this area leads to the finding, that the design, prototyping, and evaluation of privacy-preserving personalized public displays may be considered a difficult task: At least eight challenges call for careful consideration. Specifically, privacy on personalized public displays is a sensitive issue. The following three sections address the research question proposed in Section 4.1 with specific scientific contributions: Section 10.1 presents a threat model (C1) that may be used to assess the privacy threats to public display systems; the model also allows to compare different systems. Section 10.2 presents three novel countermeasures (C2) to privacy threats, i.e., visual multiplexing, visual highlighting, and visual interaction. To address the eight challenges and to consolidate C1 and C2 in an integrated process, Section 10.3 proposes a new methodology to design, prototype, and evaluate privacy-preserving personalized public displays (C3). Overall, C1–C3 strive to provide a holistic approach to the complex issue of designing privacy-preserving personalized public display systems.

10.1. Privacy Threat Model

As mentioned in Section 7.2, there appears to be no universal threat model applicable to public displays. The general concept of threat models, however, is well established in engineering and research, see Sections 6.5 and 7.5. It thus seemed reasonable to use an approved concept as the base of a privacy threat model for personalized public display systems. Subsection 10.1.1 explains the selection of an appropriate theoretical grounding in depth. Afterwards, Subsection 10.1.2 presents details about the derived design of the privacy threat model.

10.1.1. Deriving a Theoretical Grounding

The Open Web Application Security Project (OWASP) periodically issues reports on the ten most critical (security) risks that web applications may be subject to. The goal of the project is to raise awareness about possible risks by providing a concrete tool, i.e., the reports, that various stakeholders, e.g., “developers, designers, architects, managers, and organizations” [227, p. 3], may use. Renowned organizations reference the OWASP reports, for example, the Federal Trade Commission (FTC) [77]. The 2013 edition of the “OWASP Top 10” [227] report contains a visualization about how attackers may exploit security weaknesses in order to take advantage of certain technical and business impacts. This visualization is shown in Figure 10.1. Together with concrete instances of threat agents, attack vectors, security weaknesses, security controls, technical impacts, and business impacts this constitutes a threat model.

The stakeholders addressed by the project (not to be confused with the stakeholders introduced in Subsection 2.2.1) may use this threat model to guide their design and development processes to harden

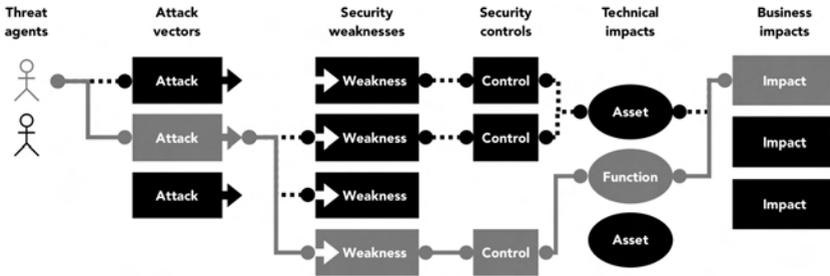


Figure 10.1.: Adapted visualization of application security risks as presented in OWASP Top 10 – 2013 [227, p. 5]. “Attackers can potentially use many different paths through your application to do harm to your business or organization. Each of these paths represents a risk that may, or may not, be serious enough to warrant attention” [227].

specific web applications. The OWASP reports propose sets of threat agents, attack vectors, and so forth. However, the threat model may be adapted to the needs of a specific project, for instance, by adding new threat agents to the list. Actually, this is what the 2013 report prompts stakeholders to do:

In the long term, we encourage you to create an application security program that is compatible with your culture and technology. These programs come in all shapes and sizes, and you should avoid attempting to do everything prescribed by some process model. Instead, leverage your organization’s existing strengths to do and measure what works for you. —*The OWASP Foundation* [227]

Though the structure of the OWASP threat model may be applicable to a privacy threat model for public display systems, some of the actual contents, for example, the attacks vectors, may be less suitable. The remainder of this subsection thus explains how these contents

were redefined. There is a special focus on the attack vectors, i.e., what attacks can be performed, as they entail the security weaknesses as well as the security controls.

Section 6.5 introduced the STRIDE threat model as proposed by Her-
nan et al. [93]. Due to its generic applicability, the STRIDE threat
model has been used successfully in other contexts [194]. To further
underpin the flexibility and applicability of the model, the STRIDE
categories can be mapped to the six security properties of confiden-
tiality, integrity, authentication, authorization, availability, and non-
repudiation [198, 246]. As the model has been successfully applied to
numerous contexts and scenarios, it may also be a good candidate to
serve as the grounding of a privacy threat model contentwise.

There are alternatives to STRIDE, for example, EBIOS¹, SP 800-30²,
Octave³, MEHARI⁴, ISO/IEC 27001⁵, the Austrian IT Security Hand-
book⁶, or the IT-Grundschutz Catalogues by the German Federal Of-
fice for Information Security [84]. However, these alternatives are ei-
ther rather complex, less specific, or not available in English. More-
over, some are commercial products, that may not be used for free.
This further underpins the rationale for choosing STRIDE as the the-
oretical grounding for a privacy threat model for public displays.

A user study was carried out to evaluate the suitability of STRIDE as
the base of a privacy threat model. The study aimed at finding answers
to RQ1 (“What are main privacy threats on public displays?”) and its

¹<http://www.ssi.gouv.fr/guide/ebios-2010-expression-des-besoins-et-identification-des-objectifs-de-securite>, accessed: July 13, 2015

²<http://csrc.nist.gov/publications/PubsSPs.html>, accessed: July 13, 2015

³<http://www.cert.org/resilience/products-services/octave>, accessed: July 13, 2015

⁴<https://www.clusif.asso.fr/en/production/mehari>, accessed: July 13, 2015

⁵http://www.iso.org/iso/home/store/catalogue_ics/catalogue_detail_ics.htm?csnumber=54534, accessed: July 13, 2015

⁶<http://www.digitales.oesterreich.gv.at/DocView.axd?CobId=23263>,
accessed: July 13, 2015

three sub-questions (“Is there a privacy threat model for public displays?”, “To what extent are existing models applicable?”, and “Which application scenario requires the most privacy?”), see Section 4.1. To further support the evaluation of these questions, the five hypotheses listed in Table 10.1 were defined.

Table 10.1.: Hypotheses used in the user study about the suitability of STRIDE as the base of a privacy threat model.

ID	Hypothesis
H1	Information disclosure constitutes a major privacy threat on public displays.
H2	Denial of service constitutes a minor privacy threat on public displays.
H3	Users accept showing personalized content on public displays automatically.
H4	Reading personal messages on public displays requires the most privacy.
H5	Playing games on public displays requires the least privacy.

The first two hypotheses are based on the assumption that revealing personal information in public is naturally perceived as undesirable (H1), while the consequences of unavailable systems are likely to be underestimated by the participants (H2). H3 is motivated by the large body of related work on this topic. H4 and H5 were defined following discussions of typical public display applications and their relative sensitivity to privacy concerns: Personal messages may obviously contain much private information [26], while games do usually not contain much user-related content, except for, e.g., high score lists.

Apparatus and Material

The study was carried out as an online survey available in English and German. To develop a successful design, it is crucial to carefully assess the user requirements, which is why an online survey was chosen. This way, the design of the threat model would be aligned to the expectations of the actual users, and may be elevated from a mere software engineering tool to an application design tool. The survey was structured in five parts as presented in Table 10.3.

Part 1 briefed the participants with a short introductory text. To avoid setting participants in a privacy-driven mindset, which could have biased the answers, the text did not put the focus of the study on “privacy” or “security,” but rather on “using digital public displays” in general. The introductory text also presented the common terms of participating in such a study. Part 2 gathered demographic data about the participants. All questions in this part were mandatory, but some questions, for example, the gender, provided the option to explicitly decline answering.

Part 3 assessed the participants’ everyday usage of technology. Each question was mapped to one or more STRIDE categories. This allowed to analyze the participants’ general attitude towards particular privacy aspects and STRIDE categories [53] and could also be useful to draw further conclusions in the evaluation. All questions in this part were mandatory. Part 4 contained questions about the participants’ use of digital public displays. There was a pair of questions for each STRIDE category, and an additional pair (P4.17–P4.18) for de-contextualization. The questions P4.19 and P4.20 asked for the users’ acceptance towards showing personalized content automatically.

The questions in a pair were phrased inversely to allow for a more precise assessment of the participants’ opinion and for the detection

of randomly filled in surveys. The questions P4.1–P4.20 had to be answered on a 5 point Likert scale, i.e., “strongly disagree,” “disagree,” “neutral,” “agree,” and “strongly agree.” Question P4.21 asked for “the biggest threat to your [the participants] privacy” as free text. Question P4.22 asked to rate the importance of privacy on a 5 point Likert scale, i.e., “very unimportant,” “unimportant,” “neutral,” “important,” and “very important,” for the ten actions on interactive public displays listed in Table 10.2. They were chosen based on the analysis on related work [5, 26, 67, 69, 180, 195, 231], see Section 7.5, and are supposed to take place in a well-frequented shopping mall. There was an example of each action to avoid misinterpretations by the participants. The questions P4.1–P4.20 and P4.22 were mandatory. Part 5 provided space for written comments and general feedback.

Table 10.2.: Actions on interactive public displays used in the user study about the suitability of STRIDE as the base of a privacy threat model.

ID	Action
A1	Reading personal messages
A2	Browsing your calendar
A3	Browsing pictures
A4	Watching videos
A5	Browsing websites
A6	Using a map
A7	Getting directions
A8	Checking traveling plans
A9	Playing games
A10	Using social networks

Participants

The survey was disseminated via social networks, e-mails, and bulletin boards. 118 participants answered all mandatory questions. There were 95 male and 23 female participants, 32.6 years of age in average ($s = 12.86$, $min = 14$, $max = 73$). 111 participants lived in Germany and 106 chose the German version of the survey. 105 participants claimed to own a smartphone; 94 of those participants strongly agreed to use their smartphone for other things than placing calls. 12 participants selected “middle education,” 44 “secondary education,” 50 “bachelor’s or master’s degree,” and 9 “doctorate.” 3 participants preferred to not disclose their level of education. All participants could take part in a draw to win one of five \$25 app store vouchers after completing the survey. Unfortunately, the amounts of male and female answers as well as ages are not evenly distributed. Based on the analysis of distribution in a histogram, four clusters (ranging from 0–37, 38–51, 52–62, and 63–99 years of age) were introduced for additional statistical analysis. Furthermore, most answers stem from regions in central Europe (mainly Germany) and the United States only. Though this might not affect the outcomes of the study negatively, it would have been interesting to gain insights from other regions as well.

Procedure

The online survey was publicly available for ten days. The order of the parts 1–5 was the same for all participants, while the order of questions P3.1–P3.13 and P4.1–P4.20 varied randomly for each participant. The questions P4.21 and P4.22 were always asked last to avoid setting participants in a privacy-driven mindset too early. Pre-tests showed that participants required 10 minutes to fill in the survey in average. Browser cookies were used to limit each participant to one trial only.

Table 10.3.: Structure of the STRIDE study. Questions in parts 3 and 4 are mapped to corresponding STRIDE categories. For the sake of brevity, the original expression “digital public display” has been replaced by the abbreviation “DPD.”

ID	Description/Question	STRIDE
Part 1	Briefing	
Part 2	Demographics (age, gender, residence, education, smartphone usage, DPD usage, ...)	
Part 3	Technology in everyday lives	
P3.1	I change some of my passwords from time to time.	S
P3.2	I use different passwords for my accounts, subscriptions etc.	S
P3.3	I use encryption for my private e-mails.	I
P3.4	I electronically sign my private e-mails.	T, R
P3.5	When entering the PIN for my credit or debit card, e.g., at an ATM, I always check that no one is looking over my shoulder.	S
P3.6	I always make sure that a website uses an encrypted connection before I enter my username or password, i.e., “https://” instead of “http://.”	S, T
P3.7	I frequently use the “incognito” or “private” mode of my browser that prevents the sites I visit from collecting information about me.	R, I, D
P3.8	I post messages, photos etc. on social networks, e.g., facebook or twitter, from time to time.	R, I
P3.9	I have an antivirus software installed on my computer.	T, D, E
P3.10	I keep my software up to date in terms of security.	T, E
P3.11	I use services such as Dropbox, iCloud, or Skydrive to store my data in the cloud.	T, I, E
P3.12	I use public WiFi networks, a.k.a. Hot Spots.	T, I, D, E
P3.13	I use public computer terminals (e.g., internet stations at airports or other public buildings) to log in my e-mail account, for example.	T, I, D, E

Table 10.3.: Structure of the STRIDE study (continued).

ID	Description/Question	STRIDE
Part 4	STRIDE on DPDs	
P4.1	I would enter my username and password on a DPD.	S
P4.2	My username and password have to be processed and stored securely.	S
P4.3	I do not mind if anyone catches what I input on a DPD.	T (input)
P4.4	Some of my input on a DPD may be confidential.	T (input)
P4.5	I'm fine with user-contributed content on a DPD.	T (output)
P4.6	I would prefer administrated content, i.e., content selected by an editorial committee, only on a DPD.	T (output)
P4.7	I would prefer the DPD to keep a record of what has been shown on the screen for a certain period of time, e.g., 1 hour, 1 day, or 1 week.	R (output)
P4.8	I would prefer the DPD not to keep a record of what content was shown when.	R (output)
P4.9	I would prefer the DPD to record who used it when and how.	R (input)
P4.10	I would prefer the DPD to not keep records on who used the system at what time and in which way.	R (input)
P4.11	While others are around, the DPD should only show content that cannot be directly related to me.	I
P4.12	The DPD should not show personal content, e.g., e-mails, text messages, calendars, or traveling plans.	I
P4.13	If the DPD would become unresponsive while I am interacting with it, I would feel uneasy.	D
P4.14	If the DPD would become unresponsive while I am interacting with it, it would not affect me much.	D
P4.15	If I could gain access to another user's content on a DPD without their permission, I would not use the system for my own content.	E
P4.16	If someone else could gain access to my content on a DPD without my permission, I would use my own content on the system.	E
P4.17	I do not mind if others see me using a DPD.	D*
P4.18	While using a DPD, I prefer not to be observed by others.	D*

Table 10.3.: Structure of the STRIDE study (continued).

ID	Description/Question	STRIDE
P4.19	Personal content, e.g., e-mails, text messages, calendars, or traveling plans should be shown automatically on a DPD.	
P4.20	A DPD should show personal content, e.g., e-mails, text messages, calendars, or traveling plans on demand.	
P4.21	While using a DPD, I think the biggest threat to my privacy would be ...	
P4.22	Imagine you were using a DPD in a well-patronized shopping mall. How important is your privacy to you while ...	
Part 5	Comments, debriefing	

Results

The answers to questions with a 5 point Likert scale were mapped to numeric values ranging from -2 (“strongly disagree” or “very unimportant”) to 2 (“strongly agree” or “very important”) with 0 as the neutral element. The overall Likert score for each STRIDE category per participant was computed by adding the Likert scores of the corresponding question pairs. Averages of the positive and negative questions are not bound to aim towards zero, since they are not designed as perfect mathematical counterparts. To the contrary, the average scores may thus provide more differentiated and reliable results. Also, averaging Likert scores is common practice in literature. Finally, it was also taken into account whether the individual question was negatively or positively phrased. For example, there were two questions related to spoofing: P4.1 and P4.2. The overall score for spoofing was thus calculated as: $Likert_S = -P4.1 + P4.2$. Consequently, the values for the overall scores range from -4 to 4, while 0 still indicates the par-

participant's neutral attitude. Summing up the overall Likert scores of all participants provide insights about the relevance of the particular STRIDE category, see Table 10.4. The (implicit) prioritization of privacy threats on public displays (from most relevant to least relevant threat) can thus be derived as: spoofing, repudiation, decontextualization, tampering, information disclosure, denial of service, and elevation of privilege (SRD*TIDE).

Table 10.4.: Total Likert scores and relevance of P4.1–P4.16 and P4.17–P4.18 (implicit prioritization).

ID		Total Likert score	Relevance
S		319	1
T	input	41	5
	output	35	6
R	input	272	2
	output	213	3
I		8	7
D		-245	8
E		-336	9
D*		48	4

52 participants provided free text answers to question P4.21, which were mapped to the STRIDE categories and decontextualization by 7 raters (R1–R7). Krippendorff's alpha was used to calculate the inter-rater reliabilities of all logical combinations (e.g., R1+R2, R1+R3, ..., R1+R2+...+R7) and to find the most meaningful ones, i.e., the combinations with the highest alpha values. The STRIDE mappings of the corresponding raters were then summed up and used to calculate the overall relevance of each category, see Table 10.5. The (explicit) prior-

itization of privacy threats on public displays (from most relevant to least relevant threat) can thus be derived as: information disclosure, spoofing, elevation of privilege, tampering, decontextualization, denial of service, and repudiation (ISETD*DR).

Table 10.5.: Ratings (R1–R7), sum, and relevance of each STRIDE category according to P4.21 (explicit prioritization).

ID	R1	R2	R3	R4	R5	R6	R7	Sum	Relevance
S	-	-	18	20	-	-	-	38	2
T	-	-	8	-	-	7	-	15	4
R	3	-	1	-	-	-	-	4	7
I	-	-	-	-	35	35	-	70	1
D	2	-	-	2	-	2	-	6	6
E	15	-	-	-	13	-	-	28	3
D*	7	7	-	-	-	-	-	14	5

Figure 10.2 visualizes both, relative implicit and relative explicit results, in a juxtaposition. The values for the implicit prioritization were calculated by dividing the corresponding total Likert scores in Table 10.4 by 472. (There were 118 participants and two 5 point Likert scales ranging from -2 to 2 for each STRIDE category, thus $472 = 2 * 2 * 118$.) The values for tampering and repudiation are based on the arithmetic means of the corresponding input and output values. The percentages for the explicit prioritization were calculated likewise by dividing the corresponding total scores in Table 10.5 by 104. (There were 2 raters classifying 52 answers, thus $104 = 2 * 52$.) The answers with regard to the general acceptance of automatically showing personalized content (P4.19–P4.20) result in a total Likert score of 102.

The 118 individual Likert scores for each action A1–A10 of question P4.22 were summed up and ranked accordingly. Table 10.6 shows the

results and the actions on public displays prioritized in regards to privacy demands. Privacy seems thus to be most important when reading personal messages (A1), browsing pictures (A3), as well as using social networks (A10).

Table 10.6.: Total Likert score and relevance of privacy in application scenarios A1–A10.

	Total Likert score	Relevance
(A1) Messages	204	1
(A3) Pictures	154	2
(A10) Social networks	153	3
(A2) Calendar	131	4
(A4) Videos	104	5
(A5) Websites	53	6
(A8) Itineraries	26	7
(A7) Directions	-39	8
(A6) Maps	-57	9
(A9) Games	-79	10

T-tests showed that the answers to P4.1–P4.20 are statistically meaningful: The means of the answers differ significantly from 0 (neutral); the same applies to the prioritization of A1–A10. Oneway ANOVAs did not confirm statistically significant influences of any demographic data on the (implicit and explicit) prioritization of STRIDE categories or on the prioritization of A1–A10 in terms of privacy demands.

However, some answers to questions about the participants' everyday usage of technology (Part 3) seem to have a significant influence on the prioritization (note that $df_2 = 116$ for oneway ANOVAs and $df_2 = 58$ for linear mixed-effects models with Satterthwaite approximations):

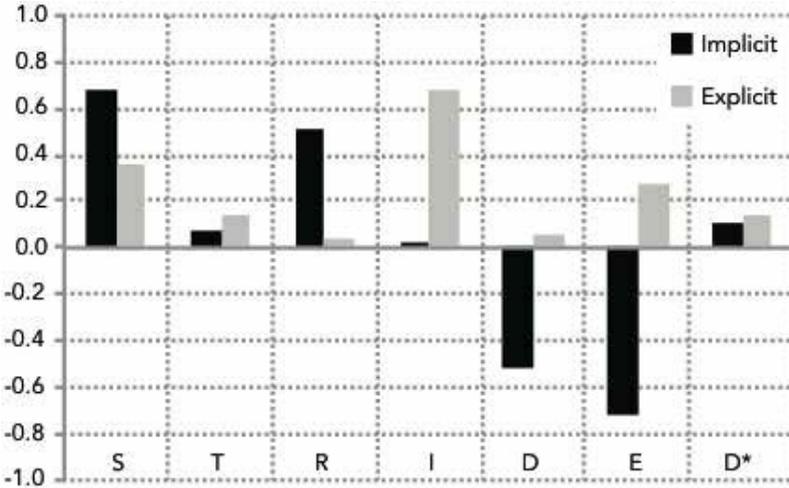


Figure 10.2.: Perceived (relative) relevance of STRIDE categories and decontextualization (D*) based on Table 10.4 (implicit results) and Table 10.5 (explicit results). 1.0 corresponds to “perceived as major privacy threat,” -1.0 corresponds to “perceived as minor privacy threat.”

Participants that rate spoofing as a major threat also tend to (i) only use encrypted websites when entering passwords ($F(1, 116) = 11.424, p = 0.001$), (ii) change passwords from time to time ($F(1, 116) = 5.373, p = 0.022$), (iii) frequently use incognito browsing ($F(1, 116) = 6.673, p = 0.011$), (iv) make sure that no one is looking over their shoulder when entering a PIN at an ATM ($F(1, 58) = 4.081, p = 0.048$), and (v) use different passwords for different accounts ($F(1, 116) = 3.911, p = 0.050$).

Participants that rate tampering as a major threat also tend to (i) avoid public WiFi networks ($F(1, 116) = 7.757, p = 0.006$), (ii) have anti-virus software installed ($F(1, 58) = 5.714, p = 0.020$), (iii) frequently use incognito browsing ($F(1, 116) = 5.113, p = 0.026$), and (iv) avoid public com-

puters for, e.g., checking their e-mails ($F(1, 116) = 3.943, p = 0.049$). Participants that rate information disclosure as a major threat also tend to (i) avoid public computers ($F(1, 116) = 14.130, p < 0.001$), (ii) avoid public WiFi networks ($F(1, 58) = 8.326, p = 0.005$), and (iii) only use encrypted websites when entering passwords ($F(1, 116) = 4.941, p = 0.028$).

Participants that rate repudiation as a major threat also tend to (i) post messages, photos etc. on social networks from time to time ($F(1, 58) = 8.839, p = 0.004$) and (ii) avoid public WiFi networks ($F(1, 58) = 6.711, p = 0.012$). Participants that rate decontextualization as a major threat also tend to avoid public computers ($F(1, 58) = 9.026, p = 0.004$).

Discussion

The results address RQ1 in two ways, i.e., explicitly and implicitly defined threats. Question P4.21 investigates the first ones by asking for the “biggest privacy threats” while using digital public displays. Questions P4.1–P4.16 examine the latter ones by asking questions related to each STRIDE category as well as decontextualization (P4.17–P4.18). Both results will be discussed in more depth below.

First of all, though the answers to question P4.21 were provided as free text, they did not reveal any other significant threats than those covered by the STRIDE categories. Seven raters were asked to identify the corresponding STRIDE categories (and decontextualization), which yielded reasonable inter-rater reliabilities, see Table 10.5. The interpretation of the explicit prioritization supports H1, since information disclosure has been declared the most significant privacy threat on public displays. The results also support H2, as denial of service is not perceived as a major privacy threat. The verification of H1 thus justifies the large body of related work presented above and further underpins the importance of further research in this domain.

The results of questions P4.1–P4.16, which tried to implicitly assess the relevance of each STRIDE category, also support H2: Denial of service is apparently not perceived as a main privacy threat on public displays, as it is ranked second to last. At the same time, however, the implicit evaluation seems to disprove H1, because participants expressed a (close to) neutral attitude towards this privacy threat. Spoofing, repudiation, decontextualization, and tampering appear to be more important to the participants. This outcome could be a result of the phrasing of the questions, which could have biased the participants. For example, a more biased version of P4.11 would be: “strangers” instead of “others;” “traced back” instead of “directly related.”

By looking at the total Likert scores, see Table 10.4, and the results of the raters’ mappings, see Table 10.5, STRIDE may be a suitable base for a privacy threat model for public displays. However, to further underpin this base, the meaning of the letter D should be changed from denial of service to decontextualization (referred to as D*), as the latter is perceived as a more prevalent privacy threat: $Likert_D = -245$ can be interpreted as “perceived as a subordinate privacy threat,” while $Likert_{D^*} = 48$ can be interpreted as “perceived as a moderate privacy threat.” There is also a significant difference between the total mapping scores in Table 10.5.

Furthermore, the ANOVAs showed that some of the participants’ self-reported technology habits seem to have a significant impact on the prioritization of corresponding STRIDE categories. For example, participants that tend to rate spoofing as a major privacy threat also tend to make sure that no one is looking over their shoulder when entering a PIN at an ATM and frequently use incognito browsing. This observation supports the premise that the definitions of these STRIDE categories match the participants’ understanding—“The STRIDE categories measure the right things.” To design interactive public displays

tailored to the privacy demands of such users, designers could, for example, equip ATMs with mirrors that help to see who is standing behind you or avoid unnecessary data gathering, e.g., “autocompleting” forms, and explicitly advertise this frugality.

The magnitude of the total Likert score to the questions P4.19–P4.20 shows that users accept personalized content on public displays: $102 \div 472 = 0.22 \gg 0$ indicates a positive attitude, which supports H3. It further emphasizes the importance of a privacy threat model for public displays to design privacy-aware systems. The answers to question P4.22 indicate which actions on public displays have the highest and which have the lowest privacy demands, see Table 10.6: Participants have a high demand for privacy while reading personal messages (A1)—which supports H4; participants have the least demand for privacy while playing games (A9)—which confirms H5.

The distribution of the remaining activities along this “dimension of privacy demands” also reveals further interesting insights: Browsing pictures (A3), e.g., looking at the latest Flickr photos or pictures that someone shared with you on Facebook, seems to require a significant amount of privacy—even more than looking at personal calendars (A2). In addition, participants desire a high degree of privacy while using social networks (A10). Possible explanations for this observation are manifold: First of all, social networking may include many types of content, e.g., personal messages, appointments, photos, or videos. Social networking could thus be regarded as a superset of the corresponding individual actions A1, A2, A3, and A4. Also, it could prove the importance of social networking activities in society nowadays.

Overall, the results to question P4.22 could be interpreted as pillars spanning a design space for privacy demands on public displays. Public display designers could use the ranking presented in Table 10.6 to assess their users’ privacy requirements and create privacy-aware sys-

tems. Designs that focus on threats and privacy in such a way could positively affect user attitude and thus the overall display usage.

Limitations

This study may have been subject to some limitations. It was carried out as an online survey to allow for an easy and widespread distribution amongst participants. However, as demographic statistics reveal, this goal could only be achieved partially. To improve the limited sample size of the study, the survey could be conducted as a pen and paper version, too. This way, even more opinions of older participants—less likely to fill in an online survey—could be reflected. Moreover, a more equal distribution among female and male participants, as well as countries and levels of education would have been desirable to further reinforce the results of the study.

To remedy possible impacts of the phrasing of questions P4.1–P4.22 different sets of questions could have been used among all participants. This could also positively affect the overall Likert scores per STRIDE category, as the additional scores would allow to assess the importance of the category more precisely and widen the scope of the resulting Likert scale. For example, one 5 point Likert scale results in a range between -2 and 2; two such scales result in a range between -4 and 4 and so forth. With regard to the applicability of inferential statistics, such an expansion could thus improve the normal distribution of the results. However, drafting multiple questions that assess the relevance of each STRIDE category evenly turned out be a challenging task, as categories, such as denial of service, appear difficult to ask for with varying wordings. Another culprit could be the mapping of the free text answers to the individual STRIDE categories. As the analysis of Krippendorff's alpha revealed, taking the mappings of

all seven raters into account simultaneously would have led to statistically unreliable results. Instead, pairs and triplets of most correlating raters were identified and used in the analysis. Training the raters more intensely beforehand could have improved outcomes.

Summary

A user study was carried out to evaluate whether STRIDE may serve as the content-related base of a privacy threat model for public displays. Regarding RQ1 and the first two sub-questions, the results indicate that STRIDE can be used to model main privacy threats. However, the meaning of the letter D should be changed from denial of service to decontextualization. Thus, the modified STRIDED* threat model should be used for future analyses of interactive public displays in terms of privacy. The study also identified the relative importance of these privacy threats. There is an apparent discrepancy between the participants' explicit and implicit prioritization of those threat categories: ISETD*DR (explicit) vs. SRD*TIDE (implicit). Yet, the results imply that public display designers should especially focus on privacy threats induced by either information disclosure or spoofing. The results also define a design space for privacy demands while performing certain actions on public displays. With regard to specific application scenarios, public display designers can use this design space to build privacy-preserving systems that align with users' privacy perceptions and needs. This addresses the third sub-question of RQ1.

10.1.2. Derived Design

The previous subsection described how the theoretical grounding of the privacy threat model for public displays was chosen: The OWASP Top 10 report of 2013 [227] provides the structure that can be populated with the contents of the STRIDED* model. The final privacy threat model for personalized public display systems is depicted in Figure 10.3. It consists of six components: *purpose objects*, *agents*, *threats*, *weaknesses*, *effects*, and countermeasures. The remainder of this subsection explains each component of the privacy threat model in depth.

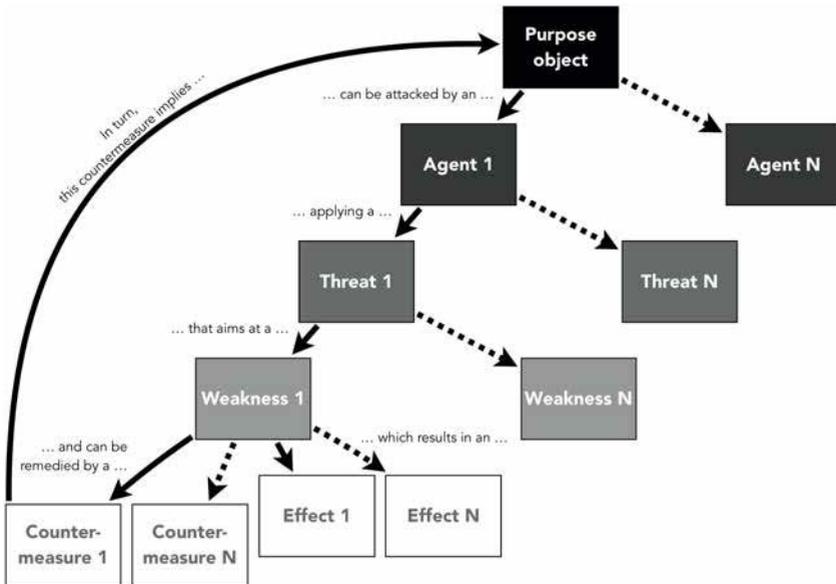


Figure 10.3.: Final design of the privacy threat model for personalized public display systems (C1).

Purpose Objects

This component is not part of the OWASP model. It has been added to the privacy threat model to allow for cyclic iterations, cf. the structure of the *composite pattern* in software engineering [80, pp. 163]. Instances of purpose objects can either be public display systems or countermeasures. In most cases, the first iteration of a threat model for a specific public display system starts with that public display system as the root purpose object. The idea is that this public display system has a certain purpose, e.g., providing users with navigational information. In subsequent iterations, countermeasures are regarded as purpose objects as they have the purpose to harden that particular public display against possible threats. These countermeasures may, however, in turn be subject to certain attacks as well. Thus, the model accounts for cyclic iterations.

Agents

The OWASP model refers to this component as threat agents. The name has been altered in the final privacy threat model to emphasize the fact that also non-malicious, i.e., neutral or benign, users may pose certain threats to public display systems. The OWASP model does not specify concrete agents, but rather provides a list of factors that can be used to classify and weigh them, see Table 10.7. The weights can be summed up to get the overall estimated “likelihood and impact level,” see Table 10.8.

For example, an attacker with “no technical skills,” who may get “low or no reward,” has “no known access” to the system, and belongs to the group of “authenticated users,” e.g., being an employee, constitutes a “low” threat, since the calculated weight computes to $1 + 1 + 0 + 6 = 8$.

Though the final privacy threat model does not require this classification of threat agents, calculating the weights may help users to concentrate on the most dangerous agents or impending threats first.

The proposed list of agents consists of the common “cast of characters” introduced by Bruce Schneier [201, pp. 23] as well as the stakeholders proposed by Alt et al., see Subsection 2.2.1. Table 10.9 lists all agents together with descriptions and examples.

Table 10.7.: Factors proposed by the OWASP model [226] to classify and weigh threat agents. Weights printed in parentheses.

Factor	Description
Skill level	“How technically skilled is this group of attackers? No technical skills (1), some technical skills (3), advanced computer user (4), network and programming skills (6), security penetration skills (9)” [226].
Motive	“How motivated is this group of attackers to find and exploit this vulnerability? Low or no reward (1), possible reward (4), high reward (9)” [226].
Opportunity	“How much opportunity does this group of attackers have to find and exploit this vulnerability? No known access (0), limited access (4), full access (9)” [226].
Size	“How large is this group of attackers? Developers (2), system administrators (2), intranet users (4), partners (5), authenticated users (6), anonymous Internet users (9)” [226].

Table 10.8.: Likelihood and impact levels of factors as proposed by the OWASP model [226].

Calculated weight	Description
0 to <3	Hight
3 to <6	Medium
6 to 9	Low

Table 10.9.: Agents of the privacy threat model for personalized public display systems.

Name	Description/Examples
Alice, Bob, ...	Benign users (with good intentions)
Marvin, Eve, ...	Malicious users (with bad intentions)
Carol, Dave, ...	Passersby or bystanders (with neutral intentions)
Space owner	Business owners, or municipalities
Display owner	Business owners or municipalities as well as companies, e.g., Wall or Ströer
Display provider	Companies, e.g., Wall, Ströer, or JCDecaux
Content provider	News agencies or TV stations

Threats

The results of the user study presented in Subsection 10.1.1 suggest to use STRIDED* as the list of threats for the privacy threat model for personalized public display systems. Additionally, to account for benign users, referred to as “Alice, Bob, ...,” the list contains an item labelled “normal usage.” Table 10.10 presents all STRIDED* threats along with descriptions and examples.

Table 10.10.: Threats of the privacy threat model for personalized public display systems.

Name	Description/Examples
Normal usage	(i) Alice uses the system with good intentions, but may still reveal private information unintentionally. (ii) Carol walks past Alice and catches some piece of information. (iii) The barkeeper watches Bob regularly check a gambling site and thus refrains from granting Bob any more credit.
Spoofing	(i) Marvin installs a fake keypad to get the users’ PIN numbers. (ii) Marvin puts a foil in front of the screen of the public display that shows false static content. (iii) Marvin pastes colored duct tape on the screen to make information disappear. (iv) Marvin spoofs a static reference image next to the display. (v) Marvin spoofs the display’s brand name (to gain trust, for example). (vi) Marvin creates the illusion of a fire next to the display (e.g., by creating smoke), which causes Alice to rush away without logging off. (vii) Marvin installs a fake surveillance camera to create the illusion of security, but uses that camera to spy on people’s interactions. (viii) Marvin installs fake public displays at a random sites (e.g., a fake ATM next to a bar).

Table 10.10.: Threats of the privacy threat model for personalized public display systems (continued).

Name	Description/Examples
Tampering	<p>(i) Marvin manipulates an ATM keypad to get the users' PIN numbers. (ii) Marvin manipulates screens so that they do not show information of specific colors (e.g., omit all red bank account numbers). (iii) Marvin injects delicate content onto Alice's screen. (iv) Marvin moves the public display to be able to spy on the user's input. (v) Marvin prevents others to approach the public display once Alice began to use it. This way Alice will not be disturbed or prevented from using the (manipulated) public display or Marvin may spy on her. (vi) Marvin fabricates an audience that influences Alice's behavior to his advantage. (vii) Marvin gains access to surveillance cameras and uses the video feeds to spy on Alice. (viii) Marvin deactivates any surveillance cameras. (ix) Marvin trains the algorithm used to distinguish public from private data with false data, so that Alice's private information will be shown in public. (x) Marvin makes Alice believe (e.g., by false static signage or a fake story) that a (manipulated) public display can be used for a specific task (e.g., online banking). Alice might thus reveal her PIN. (xi) Carol is a space owner. He is annoyed by the loud keyboard attached to a public display and thus changes the keyboard. (xii) Carol is a space owner. He decides to make changes to an output device (e.g., the screen) so that it better suits his needs (e.g., more appealing). (xiii) Carol is a space owner. He moves the public display for some reason. As a result, the screen can now be observed in a reflection (e.g., mirror or window).</p>

Table 10.10.: Threats of the privacy threat model for personalized public display systems (continued).

Name	Description/Examples
(Non-)Repudiation	(i) Alice did not read Bob's message, but the audit trail claims that. (ii) A faulty algorithm withholds Bob's message from Alice; a faulty algorithm creates a wrong connection between Bob's message and Alice. (iii) Due to bad UI design, Alice overlooks Bob's message. (iv) Marvin cannot be blamed for any type of attack due to a lack of evidence or the purposeful removal of such.
Information disclosure	(i) Due to a bad or unsuited input device, output device, or UI, Alice unintentionally shows private information in public. (ii) Marvin exploits unsuited input/output devices to gather sensitive information (e.g., a PIN entered via gestures or a screen that is too large). (iii) Marvin spies on Alice's screen content, e.g., via shoulder surfing or surveillance cameras. (iv) By using social engineering, Marvin tricks Alice to provide some (private/sensitive) data. (v) Marvin spies on Alice's input, e.g., via shoulder surfing or surveillance cameras. (vi) Marvin observes Alice using an ATM and thus concludes that she's got some cash afterwards. (vii) Alice, Marvin, Carol, or the content provider use keywords (accidentally/intentionally) so that an algorithm classifies a particular content as public though it contains (mostly) private/sensitive information.

Table 10.10.: Threats of the privacy threat model for personalized public display systems (continued).

Name	Description/Examples
Denial of service	(i) Marvin causes the public display to freeze, which may result in a prolonged display/exposure of Alice's private data. (ii) Marvin or Carol block physical access to input/output devices. (iii) Marvin overloads the public display (e.g., via external user input), which in turn does not respond to Alice's actions (e.g., log out) anymore. (iv) The display provider puts the public display in maintenance mode to install updates that fix a broken algorithm or UIs; the public display may then become unavailable/unresponsive.
Elevation of privilege	(i) Alice grants (more) rights to Bob/Carol/Marvin by accident due to a badly designed input device, automatic algorithm, or UI. (ii) Marvin or Carol may control Alice's session from afar due to an inappropriate input device. (iii) Marvin or Carol may read more than allowed to due to an inappropriate output device (e.g., screen that is too large). (iv) Marvin pretends to be a passerby first. Once Alice logged in the public display, Marvin takes over her session forcefully. (v) By using social engineering, Marvin tricks Alice to give him more rights.

Table 10.10.: Threats of the privacy threat model for personalized public display systems (continued).

Name	Description/Examples
Decontextualization	<p>(i) Alice or Carol observe Bob using a secure input device. They thus falsely conclude that Bob hides something. (ii) Alice or Carol draw conclusions based on shown content out of context. (iii) Alice or Carol see Bob using a public display in front of a travel agency. They thus conclude that Bob is about to go on a trip. (iv) Alice or Carol see Bob using a public display while Bob is surrounded by a group of soccer fans. They thus conclude that Bob belongs to that group. (v) Alice or Carol draw wrong conclusions by observing Bob using a specific public display (e.g., using an ATM not to get cash but to recharge a prepaid phone card). (vi) Bob is unaware of the context of his own information due to unsuited output devices or unsuited content (e.g., not all content is shown on a small screen). (vii) While Alice reads something on a public display, the content provider changes the content to, e.g., adult contents, possibly by accident. (viii) Alice, Marvin, Carol, or the content provider use keywords (accidentally/intentionally) so that an algorithm classifies a particular content as public though it contains (mostly) private/sensitive information.</p>

Weaknesses

The list of weaknesses has been compiled based on the extensive reading of 120 related publications, see Subsection 7.5.1. Additionally, some of the challenges introduced in Chapter 9 constitute weaknesses that may be attacked by the previously introduced agents or threats. Table 10.11 lists all weaknesses together with descriptions and examples.

Countermeasures

Based on the results of the extensive literature survey (see Subsection 7.5.1), the privacy threat model for personalized public displays contains the countermeasures listed in Table 10.12. In addition to the descriptions and examples provided here, Subsection 7.5.1 also elaborates each countermeasure in detail, see pp. 106.

Effects

The technical impacts proposed by the OWASP model were consolidated with the business impacts [226]. The distinction has been removed as it would have rendered the final privacy threat model for personalized public displays too technical and less concrete, as the OWASP model does not provide specific examples for technical impacts. The resulting list of effects is presented in Table 10.13, along with descriptions and examples.

Relations

The privacy threat model for personalized public displays as presented up to this point is complete and useable per se. Yet, it might be too

abstract for some users in certain application scenarios. Designers of a concrete public display system might feel overwhelmed, for example, by the available options such as the weaknesses or countermeasures. The privacy threat model also contains individual relations between its components: To support and simplify the design process, a relation constrains the list of available options. Benign users, for example, may be less likely to purposely perform an attack based on tampering. Thus, the corresponding threat of “tampering” would be unavailable for the agents “Alice, Bob, . . . ,” cf. Figure 10.4.

These relations are based on the literature review, see Subsection 7.5.1, in particular Table 7.3. For instance, the first row (“Adaptive User Profiles” [6]) indicates that the countermeasure No. 20 (“logging/audit”) may be useful for public displays used in an application scenario AS-A2 (“prospective reminding”) to mitigate the threat of tampering. Evaluating and combining these implications for the remaining projects resulted in the list of relations presented in the appendix, see pp. 444. Figure 10.4 visualizes parts of the complex structure induced by these relations. As the complete visualization could not be printed in this thesis in a reasonable manner—it would be scattered across too many pages—, Figure 10.4 only shows the first purpose object (a public display), the agents, threats, and weaknesses; the effects and countermeasures are omitted. As mentioned above, these relations are a useful but not an indispensable part of the privacy threat model. Obeying the restrictions implied by the relations is not mandatory; the threat model may thus be used without the relations. However, they may be quite valuable with regard to the prototype presented in Section 11.1.

Table 10.11.: Weaknesses of the privacy threat model for personalized public display systems.

Name	Description/Examples
Input device	Keyboard, keypad, camera, touch-screen, or motion tracking
Output device	Visual (screens or signs), acoustic (alarms or sounds), tactile (vibration or cues)
Content/Information	Public, personal, private, or sensitive data
Fixed environmental factors	Location, orientation, walls, buildings, signs, or vicinity (see Section 9.3)
Dynamic environmental factors	Audience, time, or weather (see Section 9.4)
Surveillance cameras	CCTV or IP cameras
Un-supervised/Un-logged public display	A public display in a secluded part of a lobby
Information delivery protocol	Not restricted to technical terms (e.g., HTTP), also includes organizational information workflow
Algorithms	Software used to choose screen content (e.g., classification between public and private content)
User interface (UI)	The (graphical) user interface consisting of widgets (e.g., buttons in a browser window)
Usage/Use case scenario	What the public display is used for, i.e., its (main) purpose (e.g., getting cash at an ATM)

Table 10.12.: Countermeasures of the privacy threat model for personalized public display systems.

Name	Description/Examples
Do nothing	Refrain from any countermeasure on purpose.
Minimize	Minimize (as in Windows OS) sensitive content.
Mask	Mask sensitive content (e.g., blacken it).
Blind out personal data	Hide all personal data, leave non-personal data on screen.
Blind out all data	Hide all data, even non-personal.
Raise awareness	(i) Raise awareness for privacy threats, e.g., by indicating that someone else is looking at the public display. (ii) Show reference images of how interfaces, e.g., keypads, should look like.
Assign/Claim/Carve	Users are able to explicitly assign, claim, or carve off screen real estate to use it exclusively during interaction.
2 nd device	Mobile devices, e.g., smartphones, other displays, or other computing systems.
Moderation	Human moderators may quickly decide whether certain contents are inappropriate (also based on context, which may be complicated for a machine or algorithm).
Anonymize user data	Users could contact each other by box-numbers rather than phone numbers or e-mail addresses, or their identity would be indicated via silhouettes rather than concrete photographs.
Let social protocols handle it	New (arriving) users would usually wait for others to finish/leave a display before starting to interact with the system themselves.

Table 10.12.: Countermeasures of the privacy threat model for personalized public display systems (continued).

Name	Description/Examples
Web-Of-Trust	(i) Use external services, e.g., OpenID, Facebook, Google, to authenticate users. (ii) Use trusted hardware (e.g., TPM or Weighted Companion Cubes). (iii) Perform hardware or software integrity checks.
Plausible deniability	Show other (unrelated) content besides the actually requested information. This way, users may employ the principle of plausible deniability (i.e., they may deny that they requested the content) to protect their privacy.
Require explicit UI actions	(i) After deleting a personal text message on a public display, the system should not automatically show the next message (as it is a common behavior of desktop e-mail clients). Instead, it should wait for the user to explicitly request the next message. (ii) Interaction, especially gestures, should be unambiguous. (iii) Provide previews of contents, e.g., thumbnails.
Explicit multi-user design	(i) Interaction concept and hardware should be designed to support multi-user, e.g., avoid touchscreens that can only handle one touch event at a time. (ii) The system should be able to react/handle varying numbers of simultaneous users, e.g., by analyzing dynamic environmental factors. (iii) Implement auto-logout functions.
Minimal data gathering design	Avoid unnecessary data aggregations (e.g., IP addresses, clear names, or locations) and offer services that process as little data as possible.

Table 10.12.: Countermeasures of the privacy threat model for personalized public display systems (continued).

Name	Description/Examples
Restrict location/ groups/users	(i) Lock the system into rooms. (ii) Allow access for staff members only, e.g., via badges or shared passwords. (iii) Allow access for individual users, e.g., by using fingerprint sensors.
Abstract presentation	Use color codes, patterns, or metaphors.
Protective casing/restrict access	(i) Lock system in (theft-proof) cabinets or boxes. (ii) Provide shielding, e.g., keypad rubbers or screen filter. (iii) Make sure systems cannot be moved etc. (iv) Build walls around systems, e.g., as done for ATMs.
Logging/audit	Meticulously record all kinds of events, e.g., interaction times or user input. This way, it might be possible to detect or monitor a privacy breach and trace back its cause as well as to detect actual user behavior in order to adapt accordingly.
Acknowledging	(i) Show previews of contents first. Users may then acknowledge their selection. (ii) Users have to interact regularly to prevent auto-logouts, i.e., they have to acknowledge that the system is still used.
De-anonymize users	To alleviate privacy threats induced by defamation etc., it may be helpful to de-anonymize users, e.g., by using signatures (PGP, MIME, etc.) or providing explicit context (show profile pictures next to comments).

Table 10.12.: Countermeasures of the privacy threat model for personalized public display systems (continued).

Name	Description/Examples
Restrict interaction	Users are only allowed to perform predefined actions, e.g., by locking-down a browser to kiosk mode.
Human supervision	(i) Regular human inspections. (ii) Camera surveillance. (iii) Hardware and software integrity checks. (iv) Monitoring changes to fixed environmental and dynamic environmental factors. (v) Do not rely solely on algorithms, e.g., for classification.
Unobtrusive interaction	In some situations, the users' privacy may be at risk, if the interaction itself can be observed by others, e.g., because using the system may be embarrassing.

Table 10.13.: Effects (of the threats) of the privacy threat model for personalized public display systems.

Name	Description/Examples
Loss of confidentiality	Private or sensitive data becomes public, includes company and personal data.
Loss of integrity	Data or information has been changed unauthorizedly, both input and output.
Loss of availability	Data, information, or services become unavailable (e.g., due to deletion or denial of service attacks).
Loss of accountability	Someone or something cannot be made responsible for something they did.
Financial damage	Loss of earnings (e.g., due to denial of service attacks) or theft (e.g., abuse of credit card information).
Reputation damage	Publishing unfavorable data or information (possibly in the name of someone else) without consent.

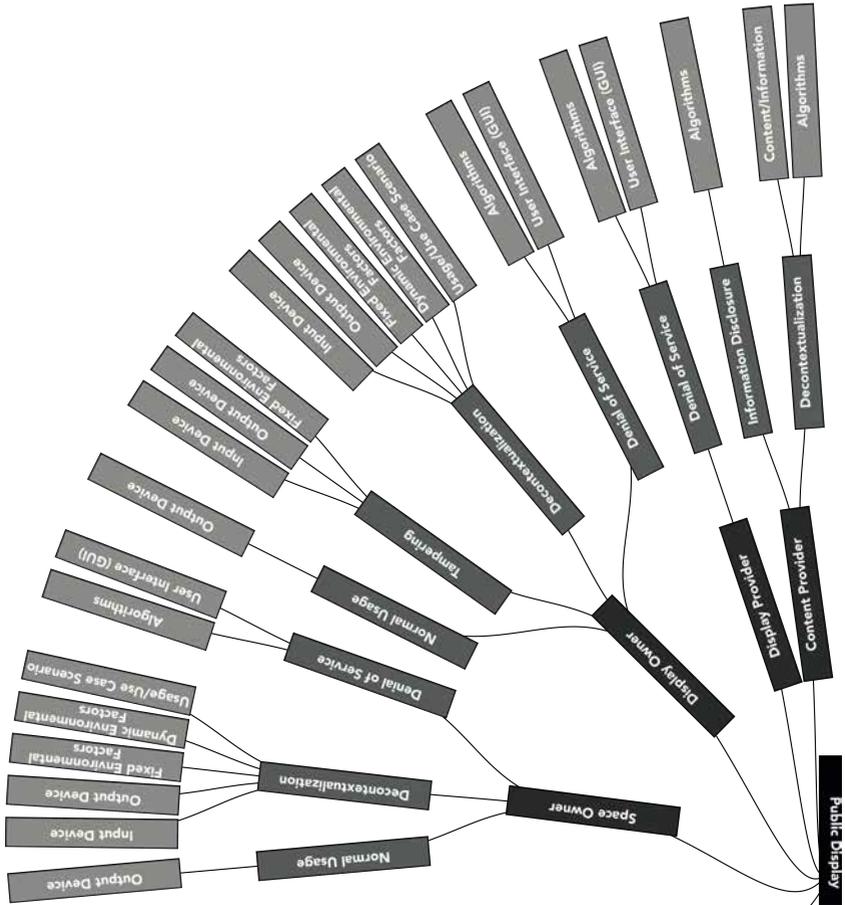


Figure 10.4.: Visualization of the relations between individual components of the privacy threat model.

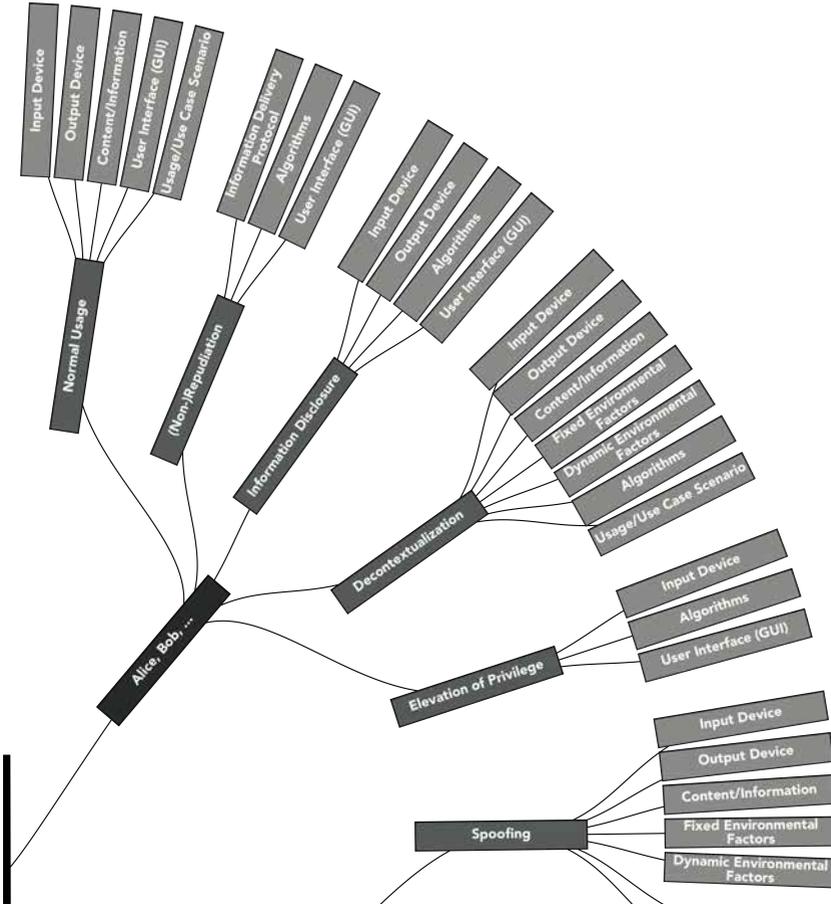


Figure 10.4.: Visualization of the relations between individual components of the privacy threat model (continued).

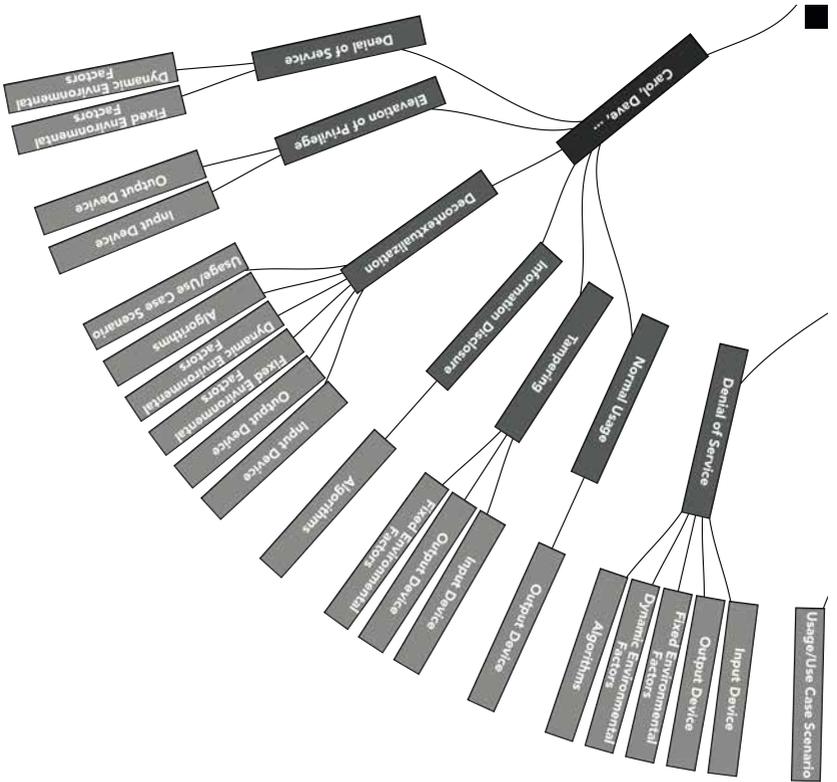


Figure 10.4.: Visualization of the relations between individual components of the privacy threat model (continued).

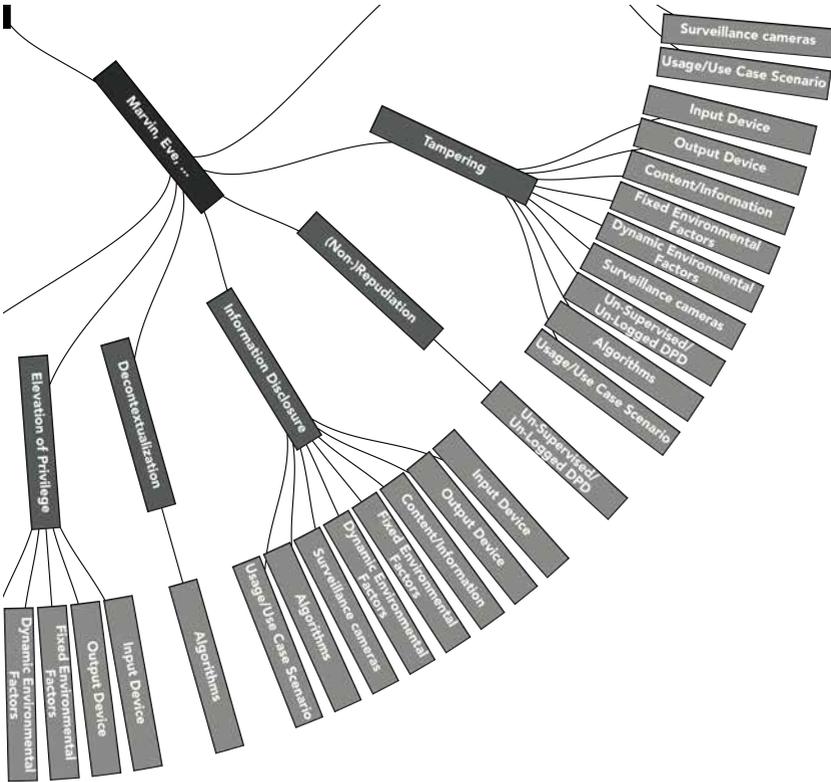


Figure 10.4.: Visualization of the relations between individual components of the privacy threat model (continued).

Examples

In order to show how the presented privacy threat model may be applied, this section reviews an example application in more detail. An ATM (purpose object) could be attacked by Marvin (agent), possibly applying a tampering attack (threat) that aims at the input devices of the system (weakness); this may result in a loss of confidentiality (effect) and could be remedied by protective casing (countermeasure).

In turn, this countermeasure (protective casing) could be attacked by Marvin (agent), possibly applying a tampering attack (threat) that aims at the input devices of the system (weakness); this may result in a loss of integrity (effect) and could be remedied by human supervision (countermeasure). In turn, this countermeasure (human supervision) could be attacked by Marvin (agent), possibly applying a denial of service attack (threat) that aims at the installed surveillance cameras (weakness); this may result in a loss of availability and could be remedied by restricting access to certain locations/groups/users (countermeasure). This scheme can be repeated until, for example, the most impending privacy threats, i.e., spoofing and information disclosure, or all STRIDED* threats have been addressed.

This cyclic iteration can possibly be continued a couple of times until no further reasonable concatenations of components come to mind. Then, all remaining combinations, e.g., agents or threats, should be considered. The underlying structure resembles something called a “tree” in graph theory. This tree may grow rapidly both in width as well as in depth, see Figure 10.4 and Section 11.1.

10.2. Countermeasures

The review of related work about personalized public display systems presented in Section 7.5 already addresses RQ2 (“What are countermeasures to those privacy threats?”) by contributing a list of 25 countermeasures. This list constitutes the major part of C2. To extend this list of countermeasures, this section presents three novel privacy-preserving approaches: visual multiplexing as described in Subsection 10.2.1, visual highlighting as explained in Subsection 10.2.2, and visual interaction as introduced in Subsection 10.2.3. All three approaches rely on personal mobile devices, i.e., smartphones, as using second devices turned out to be the most frequently used countermeasure (No. 8), cf. Table 7.10.

Personalization has been identified to be a promising approach towards display blindness, see Section 3.1. Subsection 2.2.2 and Section 7.3 point out that personalizing public displays always requires some sort of interaction. Interaction, in turn, is usually understood as a reciprocal action or influence between two parties. The ACM SIGCHI curricula for human-computer interaction [2] contains an illustration that visualizes this bidirectional relationship, see Figure 10.5.

The first two countermeasures focus on the interaction directed from the public display towards the user: Visual multiplexing allows for transferring multiple pieces of information, e.g., images, from public displays to smartphones solely based on optical communication. Similarly, visual highlighting helps users to identify relevant pieces of information on a densely populated public display; it also allows users to interact with each other, for example, by pointing to specific screen areas. The third countermeasure, i.e., visual interaction, eventually addresses the interaction directed from the user to the public display. The approach supports adaptive user interfaces that can be

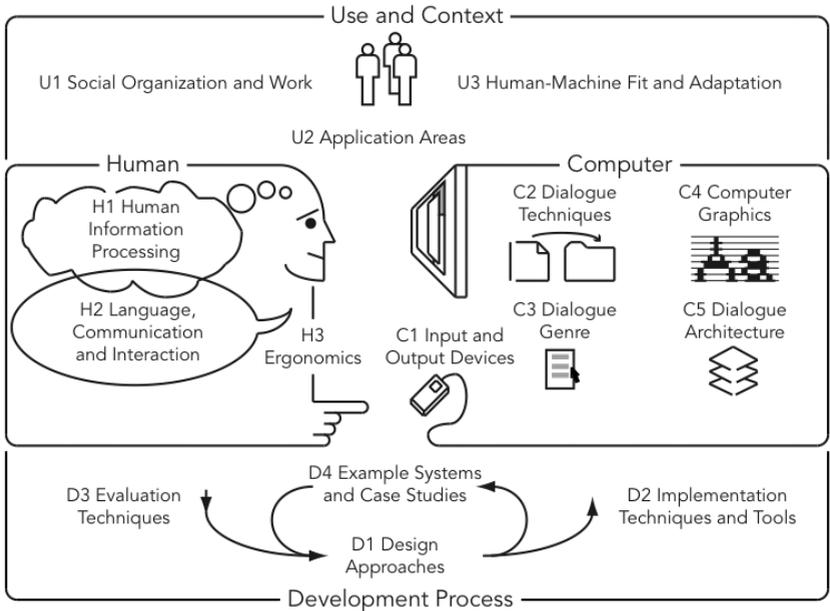


Figure 10.5.: The ACM Special Interest Group for Computer-Human-Interaction (SIGCHI) identifies sixteen topics that outline the scope their concerns. This figure visualizes some of these topics as well as their interrelationships, and indicates the bidirectional nature of interaction between humans and computers: Humans process the information shown in dialogs on computer screens, while the computer reacts to the user's input. This figure has been adapted from [2, p. 16].

tailored to suit a particular application scenario. The remainder of this section elaborates on each countermeasure in more detail.

10.2.1. Visual Multiplexing

This approach is inspired by the notion of multiplexing in electrical engineering. In this domain, multiple information signals, e.g., radio signals, are transferred over a single carrier medium, such as the air (electro magnetic waves, actually). The receiver can access individual information signals separately by demultiplexing the compound signal. Visual multiplexing applies this scheme to public displays (as senders) and mobile devices (as receivers) using visual information transfer only. The concept is based on two components: (i) A software multiplexer prepares the content, referred to as *input images* below, for a conventional public display; (ii) a mobile application dissects the components of the multiplexed information on the public display. Users can use the demultiplexer to select which information to show—they select their preferred *information channel*. Visual multiplexing may thus also help to reduce the information overload problem, see Chapter 1. The approach presented in this subsection comprises three multiplexing methods: frequency- (FDM), code- (CDM), and time-division multiplexing (TDM), cf. [102]. The difference between them is the dimension used to perform the actual multiplexing.

Visual multiplexing on public displays and visual data transfer can offer some advantages over other means of personalization, e.g., augmented reality: (i) The public display and the user's demultiplexer are not connected via conventional data channels, such as WiFi, Bluetooth, or 3G. This prevents additional data fees, especially roaming charges, and avoids any network setup, e.g., a Bluetooth pairing process. (ii) Since there is no conventional data connection, the commu-

nication between the public display and the demultiplexer is not easily traceable. This better preserves the user's privacy when using the system compared to network connections. (iii) Public displays may serve as highly visible indicators that specific information is available at a particular location, overcoming the discoverability problem of other technologies. (iv) Public displays can cater for people with and without multiplexing devices, thus providing support for a wider audience.

The remainder of this subsection presents related work on visual data transfer and visual multiplexing of public displays. Afterwards, each multiplexing method is explained in more detail. Subsection 11.2.1 presents prototypic implementations for each multiplexing method. Finally, a user study was conducted to evaluate this countermeasure and analyze the suitability of each multiplexing method with regard to contents, feasibility, and workload. Subsection 12.2.1 reports on the results of this user study.

Visual Data Transfer

QR codes are a common approach to transfer data visually to smartphones—usually links to websites, i.e., URLs. Users may scan the QR code and then navigate to the referenced website in order to retrieve related information. QR codes are well developed, broadly available, and very robust. QR codes are two-dimensional matrix codes, comparable to Aztec Codes or Data Matrix. They are an improvement to one-dimensional bar codes (e.g., EAN or UPC) that can only contain a very limited amount of data, i.e., 12 digits in case of UPC. In contrast to wireless technologies such as RFID or NFC, QR codes can be scanned from a larger distance, given that the QR code is displayed in a sufficient size. Hence, it is also possible to scan the codes in situations in

which users cannot reach the public display. Such scenarios could include the separation of users and displays by barriers, e.g., shopping windows, or situations in which additional hardware, i.e., RFID and NFC readers, cannot be installed to the display (e.g., retail TV sets).

There are 40 different versions, called levels, in the current specification for QR codes. Each version specifies how many modules a code may contain. A module, in turn, is one of the many black squares inside the QR code. In addition to the large number of specifications, there are four levels of error correction. At the highest error correction level, a QR code can still be read if 30% of its modules are corrupted. On the other hand, however, higher error correction levels reduce the overall amount of data that can be transferred in order to maintain the required information redundancy. The lowest error correction level, in contrast, maximizes the data capacity, while only 7% of the tag can be corrupted. The smallest QR code, a version 1 tag, is made up of 21 x 21 modules and can contain 25 alphanumeric characters at the lowest error correction level. The biggest QR code, a version 40 tag, can contain 4,296 alphanumeric characters at the same error correction level.

The presented approach to visual data transfer is thus based on QR codes, since they provide a comparably large data capacity and robustness. The latter is important because public displays may be installed outdoors and may thus be exposed to unfavorable lighting conditions or occlusion. Moreover, QR codes have become popular in recent years: They are a common sight on posters, billboards, or product wrappings. It can be assumed that people are aware of the concept of scanning these codes with their smartphones. Visual multiplexing as presented here requires a very lightweight infrastructure. Neither the public display nor the user's mobile device need to be equipped with dedicated hardware. The data transfer between the public dis-

play and the mobile device relies on visual means only. A special QR code is shown on a public display, which usually does not require any additional processing power. Since the data is transferred purely visually by scanning the QR code, bandwidth is not an issue and hence there is no limit on the number of parallel users. The only requirement is the visibility of the QR code. Thus, the approach scales well for large numbers of parallel users. A dedicated smartphone application is needed to scan and decode the QR code. This application could either be deployed separately or as an extension of commonly available QR code apps. This way, one may use a single application for all kinds of QR codes, including the one specified in this thesis.

Visual communication is also often used in research to transfer data between public displays and personal mobile devices: Collomosse and Kindberg [58] proposed *Screen Codes*. Their system manipulates the brightness of the display content in a regular pattern. The resulting image looks like a big QR tag superimposed on the actual content. The generated pattern may change over time, thus allowing for a constant data stream. This system is very similar to the CDM approach presented below. However, rather than enriching the display content with additional information (e.g., texts or primitive shapes), the CDM approach actually transfers the entire information itself.

In contrast to Screen Codes, *C-Blink* [147] is a system that can be used to transfer information from a user's mobile device to a public display. The screen of the mobile device emits a specific color sequence that is captured and tracked by a camera attached to the public display. This allows the user to perform certain actions, such as pressing a button or dragging an object. *FlashLight* [94] follows a similar approach. It allows for bi-directional visual communication between mobile devices and interactive tabletops. FlashLight uses RGB colors to transmit information from the tabletop to the mobile device and the mobile's

built-in flash light for the opposite direction. This approach, however, requires users to put their mobile devices directly on top of the display surface, which limits its applicability to public displays.

It is also possible to combine visual and wireless communication. *SnapAndGrab* [137] is an example system that realizes such an approach: Users snap an image of content being shown on a public display, transmit it wirelessly to the display, and then receive further information on the content being photographed, which is again transmitted wirelessly. A disadvantage of this approach is the need for the wireless link, which may incur delays and also transmission costs.

Another way to embed information in images is *watermarking* [158]. However, as Kamijo et al. [110] state, digital image watermarks may have certain limits, e.g., small data capacities or limited error correction handling. Conventional methods may in particular impact the overall readability negatively. Kamijo et al. thus propose to use invisible ink. However, this ink requires special hardware, i.e., backlight cameras, and can only be applied to analogue print media. The system presented by Yamada and Kamitani [245] can also be used without special hardware. In comparison to visual multiplexing, however, it is restricted to small payloads (64 bits) and short distances between the camera and the image.

There are also numerous radio-based communication technologies, e.g., WiFi, Bluetooth, or NFC. Yet, some may have certain disadvantages to them, for example: The senders and receivers need to be turned on all the time, which may have a negative impact on the device's battery. Moreover, the configuration and maintenance of such technologies may be difficult to some users, while taking a picture is probably a well-known task for most users.

Multiplexing Public Displays

Most related work on multiplexing public displays focusses on time- and space-division multiplexing (SDM), i.e., using the spatial position of receivers to realize SDM [95] based approaches. Kim et al. [114] proposed a generic dual-view application for a wide range of application scenarios. Their approach exploits the characteristics of Twisted Nematic (TN) LCDs and allows for SDM in its strict sense. Unlike the approach of visual multiplexing presented in this thesis, their system does not require a dedicated demultiplexer, as the demultiplexing is done by varying the user's viewing angle. However, their approach only works on TN displays and does not allow for more than two parallel information transfers. Additionally, their system has to be calibrated prior to use and may compel the user to remain in a sweet spot to optimize the viewing experience. This sweet spot is a common downside to many SDM based approaches, since the user may not be able to change the viewing angle freely in some situations. This may lead to crowding if many want to perceive the same information.

The system presented by Matusik et al. [135] follows a similar approach to show two different user interfaces at a time. Users can select the shown interface by changing their viewing angle. The authors suggest to use this system to provide two distinct views, e.g., a regular and a zoomed view for imaging software, or a normal and an outline view for web browsers. Their system could also be used to let the user toggle between different information layers, e.g., map details, by simply moving their head. Beyer et al. [27] examined the audience's behavior around cylindrical screens. Such screens also perform SDM, as the user selects the content by looking at the screen from different angles. Their work focuses on the influence of the display shape on the user and the user experience.

Alt et al. [10] also use the similar term “space multiplexing” for their approach to subdividing the display area of a public screen. Their definition of SDM, however, differs considerably from the definition used in information theory and in this thesis.

Sakurai et al. [192] focus on collaborative work on interactive multi-touch tables with an emphasis on concealing information from individual users. Their approach is based on polarization-division multiplexing, which requires users to wear special glasses. Their table top display uses motors and projector-polarizers to adjust the projected images according to the users’ positions. As a result, the system requires a sophisticated setup less feasible for use in public displays.

Olivier et al. [167] also use the term multiplexing to describe their systems called CrossBoard and CrossFlow. They emphasize the importance to balance the visibility of a public display and the ease of accessing the shown information. CrossBoard and CrossFlow thus use cross-modal attention to highlight specific information shown on a public display using TDM. The public display highlights distinct pieces of information periodically, and each time the information relevant for a user is highlighted, their smartphone vibrates. CrossFlow uses the same approach in the context of an ambient navigation system that projects a pattern of moving objects onto a surface. The Rotating Compass [189] is a very similar system to support pedestrian navigation: Lights mounted as a circle on a floor board light up in a cyclic manner. Every time the light pointing in the direction the user is supposed to go lights up, the user’s smartphone vibrates. Thus, the Rotating Compass also uses TDM to transfer individual information.

Alt et al. [10] also discuss TDM as an option. In contrast to the above mentioned systems, the TDM approach presented in this thesis is designed to operate at significantly higher speeds. While there already is a substantial body of research on visual multiplexing on public dis-

plays following the SDM and—to a lesser degree—the TDM approach, this is not the case for FDM and CDM. Since the conventional (spatial) handling of screen real estate is also well covered, visual multiplexing focusses on the less examined FDM, CDM, and TDM. In addition, a direct comparison for these methods is still missing, which is a further contribution of this subsection.

Proposed Approach to Frequency-Division Multiplexing (FDM)

FDM [102] is used, e.g., in analogue radio broadcastings. In that domain, a low-frequency signal (e.g., music ranging from 20–20,000 Hz) is modulated onto a high-frequency carrier signal (e.g., FM radio ranging from 88–108 MHz). Visual multiplexing does not directly adopt this procedure, as both signals lie in the frequency range of visible light (385–789 THz). Instead, the following strategy is proposed: Each input image is converted to grayscale. The results are then tinted red, green, and blue. The three tinted images are combined into one image using additive color mixing. The final result is displayed on the public display. The demultiplexer only shows pixels that correspond to the chosen information channel, i.e., pixels with a color value matching a defined pattern. For example, if the user selects the red information channel, only pixels with more than 50% of red will be shown.

This method supports three input images. Figure 10.6a shows a multiplexed image created with FDM. The demultiplexed results are shown in Figures 10.6b–d. The FDM method can perform multiplexing and demultiplexing in real time (i.e., at 30 fps on an iPad 3, see Subsection 11.2.1), but the original color of the input images is lost. The number of parallel information channels is limited to three due to the number of basis vectors in the RGB color space [94, 169].

Areas of public displays that have been visually multiplexed using FDM may have a distinct aesthetic look, cf. Figure 10.6. The apparent random color overlays may appear artistic and thus entice the viewer's perception of visual multiplexed public displays.

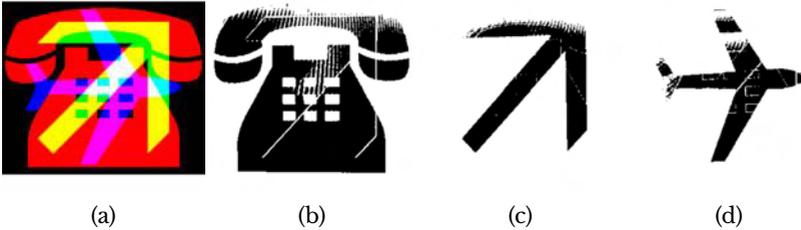


Figure 10.6.: FDM example. (a) Multiplexed image; (b)–(d) demultiplexed information channels (red, green, and blue from left to right).

Proposed Approach to Code-Division Multiplexing (CDM)

In contrast to FDM and TDM, CDM [102] does not use different frequencies or time slots to transmit each information channel. Instead, each information channel is encoded in a distinct way. This is similar to people talking in different languages. Each language represents an independent code, that can only be understood by people talking in the same language. Though the approach presented here encodes every information channel with the same code, rather than using distinct codes, CDM may still be regarded as an appropriate description of this multiplexing method.

The CDM approach used in the user study (see Subsection 12.2.1) employs QR codes and QR tags in the following way: Each input image is converted into grayscale, and scaled and cropped to 128 x 128 pixels

to reduce the data volume. Then, it is transposed from time-domain into frequency-domain by applying a Fast Fourier Transform (FFT) on every row. The upper halves of the resulting spectrums are cut off to further reduce the data volume. Subsequently, the remaining spectrums are normalized to fit a specific range. This range is dynamically computed for every set of input images in order to optimize the image quality while exploiting the maximum data capacity of QR tags. The result is compressed using `bzip2` and encoded with `base64`. The encoding avoids bytes, which cannot be represented as characters; those characters would conflict with the QR specifications.

Next, all `base64` encoded strings of all input images are concatenated with a special delimiter. The concatenated string is used to generate a QR tag, as shown in Figure 10.7a. The demultiplexer scans the QR tag, separates the information channels using the delimiter, picks out the one selected by the user, and performs the aforementioned steps in reverse order—here, a Fourier Synthesis is performed instead of an Fast Fourier Transform. Since the dropped frequencies of the spectrums cannot be recovered, they are not taken into consideration when reconstructing the image with Fourier synthesis.

While this method supports any number of input images, the current prototype (see Subsection 11.2.1) works best with three images as the current QR specifications limit the amount of data to 4,296 bytes. Figure 10.7a shows a multiplexed image created with CDM. The demultiplexed results are shown in Figures 10.7b–d. One advantage of this method is that it separates the individual information channels precisely. FDM and TDM may lack this precision due to non-optimal camera optics, lighting conditions, or timing issues. Figure 10.6b illustrates this aspect, as parts of Figure 10.6c shine through. A disadvantage of CDM is the processing time, since scanning and processing the QR tag may take up a couple of seconds. In addition, the colors

of each input image are lost. Finally, the resulting images may appear blurry, due to the spectrum cut off. Based on a constructive review comment, the CDM approach could be improved significantly:

The encoding and decoding was described in enough detail to replicate such a system—which I personally liked a lot. Nevertheless, there were quite some confusing steps [...]. First and foremost, a 128 x 128 JPEG image can be compressed to be less than 4 kB without the need to use a rather complex FFT/DCT. Particularly, because JPEG already makes use of a DCT and a subsequent LZW compression, I wondered whether the introduced encoding and decoding steps actually would lead to better results, and I have my doubts with that. Did the authors consider just sampling down the image to 4 kB and encode that one into a QR code?

—*Anonymous Reviewer at CHI 2013*

Apparently, the results shown in Figure 10.8 have a higher visual fidelity. Most importantly, the JPEG compression allowed for colored input images to be demultiplexed accordingly, see Figures 10.8e and 10.8f. Unfortunately, though, this enhancement was considered after the user study (see Subsection 12.2.1) had been carried out. The results of the study thus reflect the performance of the original approach based on FFT.

Proposed Approach to Time-Division Multiplexing (TDM)

Many current public displays already use TDM [102]: Cyclic program loops are used to display blocks of information sequentially. The users see different content depending on when they look at the screen. The TDM approach significantly increases the speed at which the information is changed by applying the following strategy: The multiplexer

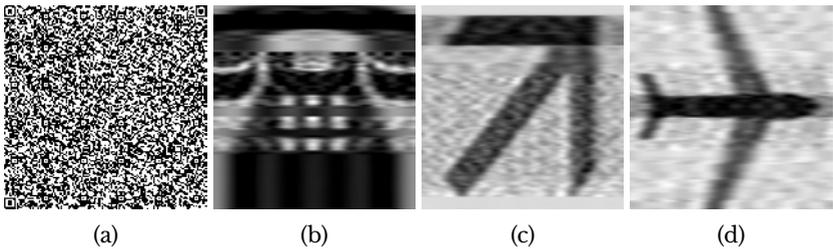


Figure 10.7.: CDM example as used in the user study. (a) Multiplexed image; (b)–(d) demultiplexed information channels.

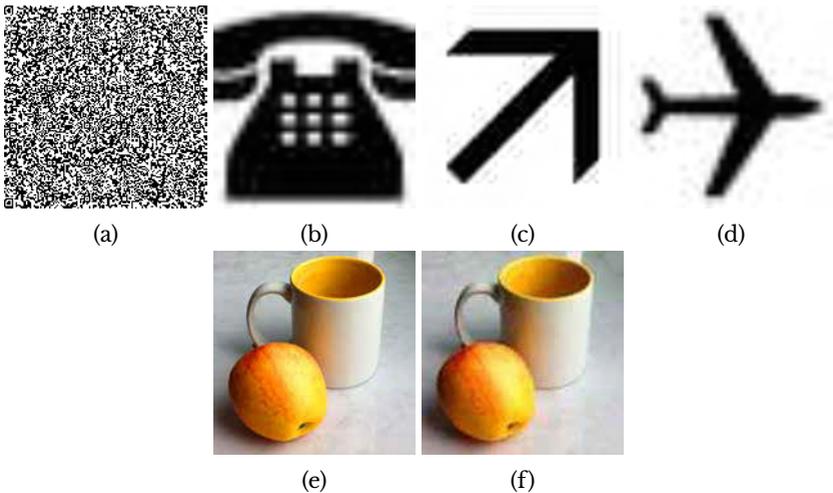


Figure 10.8.: CDM example with JPEG algorithm. (a) Multiplexed image; (b)–(d) demultiplexed information channels; (e) colored input image; (f) demultiplexed colored image.

produces a video file with a fixed frame rate, e.g., 15 fps. Depending on the frame rate, each input image is assigned a time slot with a certain length. The prototypical implementation (see Subsection 11.2.1) uses 15 fps, and thus the length of each time slot is $1 \div 15 = 0.0\bar{6}$ seconds. The first input image is used as the first video frame, the second input image is used as the second video frame, and so on. After the last input image, a special synchronization frame (depicted as a red square in Figure 10.9a) is added to the video. The resulting video is looped continuously. The demultiplexer only shows each n^{th} video frame, depending on the information channel selected by the user. The synchronization frame enables the demultiplexer to properly detect the beginning respectively the end of a cycle. For example, if the user picked information channel one out of four, the demultiplexer would only show the 1st, 6th, 11th, 16th, ... video frame.

While this method supports any number of input images, the current prototype (see Subsection 11.2.1) works best with three images due to some technical limitations as explained below. Figure 10.9a visualizes how display content changes over time. The demultiplexed images are shown in Figures 10.9b–d. The advantage of this method over FDM or CDM is that it does not manipulate the original input images, and thus retains their full color depth and resolution. Furthermore, it is significantly faster than CDM and could thus be used to visual multiplex multimedia content, e.g., videos. A disadvantage could be the noticeable flickering on the public display, as it may distract passersby [242]. Timing disparities between the public display and the mobile device could also cause *tearing*: One part of the demultiplexed information channel may contain parts of the previous or next channel, depending on whether the mobile device is ahead or behind time. This can be countered by synchronization strategies, e.g., by estimating the temporal middle of a frame.

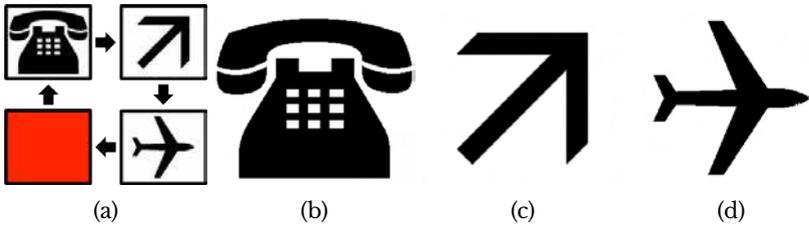


Figure 10.9.: TDM example. (a) Multiplexed image (cyclic transition over time); (b)–(d) demultiplexed information channels.

10.2.2. Visual Highlighting

One of the main application scenarios of pervasive displays is to relay information to a—possibly large—number of people. Typical use cases include conference information displays, digital signage to support navigation, as well as flight departure and arrival boards at airports. In some scenarios, the target audience of such displays will be under time pressure. At an airport, for example, some travelers may be late and thus in a rush to get to their departure gate. Flight departure boards frequently show a very large number of flights and consequently, finding the required information can be difficult, error prone, and time consuming for a traveller. Similarly, in densely crowded areas with a high throughput of people, such as entrance halls at large conferences, it may not be desirable to linger for long, and therefore the information being displayed on a public display should be spotted quickly. In such situations, it would be very beneficial if the user's attention could be efficiently directed to the one item amongst many that is relevant to that individual person.

In another application scenario, a user might want to direct another person's attention to a specific piece of information shown on a public display. One approach could be to roughly point at the corresponding

screen area; another way could be to describe the particular information or part of the display. In either case, it may not be guaranteed that the other person refers to the same piece of information at the end. Moreover, other people, e.g., bystanders or passersby, may also see the pointing gestures or overhear the oral description. This may not be desirable with regard to the user's privacy. It would thus be beneficial to direct someone's attention to a specific spot in a precise and discrete manner. Figure 10.10 shows a simple, analogous approach to visual highlighting: a "guckrohr" (a German word which roughly translates to a "looking pipe"), which is often found in zoos or parks to direct the visitor's focal attention towards, e.g., bird-nests.

In the past, a number of approaches have been proposed to address this demand for personalized highlighting of information. Frequently, they have been designed for public displays that are meant to serve large amounts of users at the same time, such as CrossFlow and CrossBoard [167] as well as the Rotating Compass [189]. While these approaches are based on crossmodal, i.e., tactile, cues emitted by personal mobile devices, visual highlighting is an alternative means that relies solely on visual cues to direct a user's attention to a specific screen item amongst many. The approach presented in this subsection uses visual markers shown on a public display to visually highlight items. The visual highlight is an overlay displayed on a personal mobile device that is pointed towards a public display. This subsection also proposes a set of criteria to classify and characterize highlighting methods for public displays and apply them to contrast the approach presented here with prior approaches.

The remainder of this subsection presents related work on previously suggested classification schemes and on existing approaches to address the issue of highlighting on public displays. Afterwards, the approach to visual highlighting as proposed in this thesis is introduced.



Figure 10.10.: A cardboard cylinder used as a “guckrohr” to realize analogous visual highlighting. The user’s focal attention is directed towards a predefined spot, here a treetop across the street. A similar concept is often used in zoos or parks to point at bird-nests, for example.

Subsection 11.2.2 presents a corresponding prototypical implementation. Next, a set of comparison criteria is proposed and motivated, while means to measure and assess the impact of each criterion on visual highlighting are outlined. These comparison criteria are used to compare and evaluate all approaches in Subsection 12.2.2. The same subsection presents a user study, which assesses the raw performance of the approach proposed in this thesis with regard to efficiency, effectiveness, and robustness.

Classification Schemes

A number of schemes to classify and compare pervasive display systems has been proposed in the past. At a fundamental level, it is possible to categorize interaction with public displays as a special case of proxemic interaction. Early work on proxemics was recently applied to large-scale displays by Marquardt et al., who proposed the Proximity Toolkit [133]. This toolkit is based on a theory of proxemic interaction and allows for rapid prototyping of applications making use of the ideas and concepts underlying proxemics.

Unlike this very generic approach, Müller et al. [151] defined a design space for interactive public displays. The focus of their framework is on forms of attraction, engagement, and interaction in the context of pervasive displays. An alternative approach was presented by Huang and Mynatt, who proposed a “design space of awareness applications categorized by the group size they are designed to support and the type of space in which they are meant to be viewed” [99]. The classification scheme proposed below complements and significantly extends prior schemes incorporating additional criteria and appropriate means to measure each dimension.

Highlighting on Public Displays

CrossFlow and CrossBoard [167] are two systems that use crossmodal cues to point out relevant information to individual users. The public display highlights distinct pieces of information periodically, and each time the user's information is highlighted, the personal mobile device signals this to the user by vibrating. This approach works with any display size and type. According to Huang and Mynatt [99], this system would suit large groups in public spaces. The two proposed use cases are an indoor navigation system (CrossFlow) and a flight departure board (CrossBoard). In these scenarios, the personal mobile device is a smartphone that has been set up and paired for a Bluetooth connection in advance. Unlike CrossFlow, which uses abstract moving patterns to indicate directions, the public display running CrossBoard remains usable by users without a personal mobile device. These users do not benefit from personalized visual highlights, though.

The Rotating Compass [189] is based on the same principle as CrossFlow. The public display is designed to resemble a compass with a needle rotating clockwise. Each time the compass needle points in the direction in which the user is supposed to move, their personal mobile device signals this through a tactile cue. This method can operate with any size and type of display and—according to Huang and Mynatt [99]—could be classified as being designed for large groups in public spaces. The main use case scenario for the Rotating Compass is an outdoor campus navigation system. The system relies on a pre-configured and synchronized smartphone, and thus cannot be used if the user does not carry a personal mobile device.

A different approach is used for the Interactive Ambient Public Displays proposed by Vogel and Balakrishnan [231], which is based on a screen equipped with a body tracking system. The public display

tracks the users' position and posture in front of the display and uses this information to select what information to show. Relevant information for an individual is shown in an user proxy bar using notification flags. Interactive Ambient Public Displays are intended to be used by pairs or small groups in private shared places, according to the design space by Huang and Mynatt [99].

Visual data transmission between public displays and personal devices is used in the Screen Codes system [58]. Images shown on a public display are overlaid with a special semi-transparent barcode similar to a QR tag. This allows for transmitting data to a personal mobile device optically while leaving the actual screen content recognizable to users. Thus, users without a personal mobile device can still use the public display, even though they cannot take advantage of the additional capabilities of the system. Even though Screen Codes do not address personalized visual highlights per se, it is mentioned as a reference for visual data transfer from a public display to a personal device. The system is designed to work with any type and size of public displays. It may be classified as suitable for installations for large groups in public spaces according to Huang and Mynatt [99].

SnapAndGrab is described by Maunder et al. [136]. It is built on a public display and a private device, both communicating via Bluetooth. This work aims to avoid the Bluetooth pairing process. To achieve this, the user takes a regular photograph of a screen area containing interesting information. They then use Bluetooth to send this picture to the public display that will accept the data transfer without any prior configuration. Once the request for additional information has been processed, additional information is sent back to the user's mobile device. This transmission does in turn not require any pre-configuration. The key drawback is, however, that the user needs to find the actual information on the public display before being able to

receiver further information. Referring to the design space by Huang and Mynatt [99], the system could be classified to handle large groups in public spaces. The public display remains usable even if the user does not have such a personal mobile device—though without the additional information.

C-Blink [147] also uses the visual channel to enable interaction with a public display. The mobile device uses its screen to emit a sequence of colors that is then captured by a camera attached to the public display. The public display processes the recorded sequence and decodes the contained information, which could be, for example, the device ID and the requested action, e.g., click, push, drag, or drop. Though C-Blink also fits in the design space by Huang and Mynatt [99] for large groups in public spaces, its basic design is less well suited for visual highlighting on public displays, and will thus not be included in the comparison presented in Subsection 12.2.2.

Proposed Approach to Visual Highlighting

The method introduced in this thesis shares the use of the optical channel with some of the systems reviewed above, but does neither require synchronization with the public display nor a wireless connection. The proposed approach to personalized visual highlighting on public displays uses a conventional public display and a personal mobile device. Similarly to Screen Codes, it relies on special code tags to transmit data through an optical channel. In contrast to Screen Codes, the code tags do not span the entire screen in a semi-transparent manner. Rather, a code tag is a composite of a regular QR tag framed by an augmented reality (AR) marker as shown in Figure 10.11. The QR tag contains information about the different visual highlights available for the current information shown on the public display.

Additionally, the QR tag is surrounded by an AR tag. This tag helps the personal mobile device to identify the position of the display in space. By applying standard AR techniques, the personal mobile device can thus correctly render additional personalized information (taken from the QR tag) onto the public display. This allows for personalized visual highlighting as shown in Figure 10.11: Users can use the screen of their personal mobile devices to look at the public display. They can then see personally relevant information that is visually highlighted over the image of the public display as it is being shown on their personal mobile device.

Azuma [18] defines augmented reality as a variation of virtual environments and provides an overview of first AR application scenarios. Early work in this area, e.g., by Caudell and Mizel [47], used see-through heads-up displays (HUD). Their system allowed to superimpose computer graphics on top of real world objects and fixate their location no matter how users moved their heads. More recently, AR has become feasible for implementation on mobile phones [19, 20, 233]. A key metaphor in this area was proposed by Bier et al. [28]: Toolglass and MagicLens. Users can look at virtual objects using special lenses as visual filters. Each lens has a specific characteristic that reveals, hides, or modifies a particular information about the scrutinized object. The approach presented here could be seen as a Toolglass or MagicLens for public displays.

Moreover, the proposed approach operates on any type and size of display and could be classified as being suitable for large groups and public spaces, according to the design space by Huang and Mynatt [99]. The public display remains usable even if the users do not have a suitable personal mobile device.

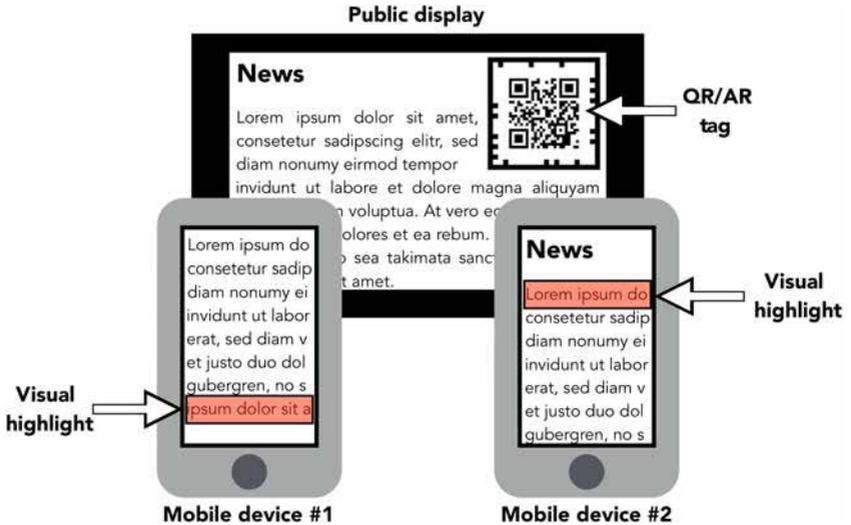


Figure 10.11.: Design of the visual highlighting approach. Two mobile devices show different visual highlights on the same public display.

Comparison Criteria

The proposed comparison criteria are motivated by the large body of existing approaches to (visual) highlighting on public displays. The following set of criteria is thus designed to be used as a universal tool for comparing different approaches for visual highlighting on public displays: Subsection 12.2.2 presents the results of a comparison of existing approaches as well as the novel method proposed in Subsection 10.2.2; the results of this comparison are discussed in Subsection 14.2.2. The comparison criteria are broken down into two sets. The first one covers the prerequisites and the second one considers the actual performance analysis. The latter one is further broken

down into four subsets, i.e., quality, quantity, reliability, and robustness. Table 10.14 provides an overview of all comparison criteria.

Prerequisites. Each visual highlighting method may use a different technical foundation. Thus, the technical prerequisites and system requirements may vary as well. This set of criteria therefore takes the following aspects into consideration: (i) the required equipment, such as certain server or mobile technologies; (ii) the effort it takes to implement the visual highlighting, concerning software as well as hardware setup; (iii) possible additional connectivity, e.g., WiFi, 3G, or Bluetooth data links; and finally, (iv) the necessary pre-use configuration, such as the pairing of Bluetooth devices. This latter is an important criterion, since it has a sizable impact on the instantaneousness and operating costs of a specific visual highlighting method. The scale to measure each criterion could be the cardinality of a list. For example, the list $L_{Equipment-A} = \{\text{computer, touchscreen, webcam}\}$ contains all three prerequisites for a specific public display system A ; thus, the cardinality of that list is $|L_{Equipment-A}| = 3$. Another public display system B could have more requirements; the cardinality of the corresponding list would thus be higher: $|L_{Equipment-B}| > |L_{Equipment-A}|$. Whether a higher or lower cardinality is “better” or “worse” likely depends on the individual situation.

Performance. The second set of the comparison criteria analyzes the performance of the individual visual highlighting method. It is broken down into four subsets, each described in detail below.

Quality. These criteria cover quality aspects of the specific visual highlighting method. Quantifying them requires the specification of ap-

propriate scales, which are shown in Table 10.14. These scales allow to assess the impact of each criterion.

Clarity defines how easily the user can identify the highlighted information. It gives information about whether the highlighting is unambiguous so that there is no doubt about what part of the public display is highlighted and thus what part of the public display is declared as relevant for the user.

Granularity defines how precisely the visual highlighting method can highlight pieces of information. A system may only be capable of highlighting rough areas, such as a whole line of text, or a system may be more precise and highlight words or even individual pixels.

Duration defines whether the method provides visual highlights at any time for an arbitrary length or whether it is bound to timing slots.

Delay defines how well the timing of visual highlights corresponds to the timing of the actual information shown on the public display. For example, there could be delays that may cause areas to be highlighted though the corresponding information has already disappeared. At worst, such a delay would lead the user to use wrong information. This criterion can be measured on a scale in seconds.

Readability defines the visual quality of the highlighted information and how the visual highlighting method may influence it positively or negatively. For example, an information printed in small letters could be zoomed in by the visual highlighting method and thus become more readable than before. In contrast, highlighting done by a colored background may lessen the visual contrast between the information and the highlight, thus decrease the readability.

Interference defines the influence of the visual highlighting method on the entire public display. For example, if the method uses special codes, like QR tags, that are shown on the public display, these codes

have an impact on the visual impression of the public display for the user. Thus, interference determines whether the system remains usable for viewers that do not use the visual highlighting capabilities. Also, visual highlights in one part of the public display may influence other information shown in different parts of the same display. Moreover, interference includes distortions between individual highlights.

Quantity. This subset covers quantity aspects for visual highlighting methods on public displays. As opposed to the previous subset, all the proposed criteria can be measured objectively on metric scales. The scalability of a visual highlighting method can also be described by these comparison criteria.

Concurrent highlighting captures how many distinct pieces of information can be highlighted in parallel. This includes any upper or lower limits and their possible causes.

Concurrent users captures how many users can be served at the same time. This includes any upper or lower limits and their possible causes. The difference between this criterion and the previous one is that though the maximum number of concurrent highlights may be limited, the method may be able to handle more users at the same time. This implies that two or more users would have to share the same visual highlight. This could be the case if two users are interested in the same piece of information, for example, at an airport while looking for their gate.

Time density captures how many sequential highlights can be shown per time unit. This criterion correlates to the notion of resolution. It encompasses the speed at which the system can operate so that the users are still able to differentiate and process the visual highlights. This includes any upper or lower limits and their possible causes.

Reliability and Robustness. The following set presents the criteria concerned with the reliability and robustness of the methods.

Availability defines how available the system is to the user with respect to conditions that have to be met so that the user may use the system. Availability is not about the current prevalence or hardware requirements of the system. For example, a system could cycle through different states, but only allow the user to synchronize in one specific state. Thus, the user may not use the visual highlighting while the system is in one of the other statuses.

Correctness defines how well the method can guarantee that the visual highlights fit the actual content. For example, the actual visible content may change while the user uses the visual highlighting. The system could continue to highlight the same region on the public display even though the corresponding content has disappeared. The user may now consider the wrong content to be relevant and pick up improper information. This criterion is, however, not concerned about the correctness of the actual visual information itself.

Environmental influences defines which external factors influence the reliability and robustness of the visual highlighting methods. For example, some visual highlighting methods require users to synchronize to the system, i.e., to look at it at the right moment for a certain period of time. If users fails to do so, they cannot use the visual highlights unless they synchronize first.

Table 10.14.: Proposed comparison criteria for visual highlighting on public displays.

Comparison criterion	Scale
(Technical) prerequisites	
Required equipment	Cardinality of a list
Effort	Cardinality of a list
Connectivity	Cardinality of a list
Pre-use configuration	Cardinality of a list

Table 10.14.: Proposed comparison criteria for visual highlighting on public displays (continued).

Comparison criterion	Scale
Performance	
Quality	
Clarity	Very ambiguous (1), somewhat ambiguous (2), not ambiguous (3)
Granularity	Block (1), words (2), pixel (3)
Duration	Occasional (1), periodic (2), continuous (3)
Delay	Seconds
Readability	Reduces readability (1), no influence (2), improves readability (3)
Interference	Interferences with conventional display use (1), interferences between highlights (2), no interferences (3)
Quantity	
Concurrent highlights	Integer
Concurrent users	Integer
Time density	Seconds
Reliability, robustness	
Availability	Occasionally available (1), periodically available (2), always available (3)
Correctness	Rarely correct (1), mostly correct (2), always correct (3)
Environmental influences	None (1), some (2), many (3)

10.2.3. Visual Interaction

The countermeasures introduced above, i.e., visual multiplexing and visual highlighting, address the communication directed from the public display towards the user. Visual interaction, as presented in this subsection, however, focusses on the inverse communication. As reasoned in Section 10.2, bi-directional communication is essential to interaction, which is—in turn—essential to personalization. As explained in Section 3.1, personalization is regarded as a promising approach to address the root of display blindness.

Interaction between users and public displays is not a new field of research. Numerous projects explored the applicability of various means of interaction to let users control such systems. However, only a few projects are concerned about the users' privacy during interaction. The related work presented in Section 7.5, Subsection 10.2.1, and Subsection 10.2.2 may already point to this fallow research opportunity; the related work presented below underlines this even more.

Section 6.2 presented arguments put forth by renowned authors that emphasize the important role of privacy as a means to avoid gradual changes in personality. In a nutshell, most people tend to adapt their behavior to the expectations and practices of the society they live in. Surveillance and data preservation are two examples of such practices that touch people's privacy. It is important to point out, that this adaptation may be a stealthy and subconscious process: Even though people claim that curtailing their privacy for whatever reason, e.g., to prevent terroristic attacks, would not impact their behavior or everyday lives at all, it often still does nonetheless.

In the context of public display systems, people might be prone to avoid interacting with such systems, since the underlying technology might be regarded as a privacy impact. Avoiding interaction implies

avoiding personalization. This runs afoul of the overall objective to address the root of display blindness. Therefore, privacy is an important aspect when designing means of interaction for public displays. Visual interaction, as proposed as a countermeasure in this thesis (C2), is an approach towards a privacy-preserving means of interaction.

Direct touch is a common type of interaction that lets users personalize the content of public displays. Besides that, interaction is also frequently realized via smartphones that wirelessly connect to the displays. While the latter approach has many benefits, for example, increased flexibility and scalability, there are also some drawbacks to it: Many smartphone-based solutions for public display interaction require radio-based connections, e.g., Bluetooth or WiFi. A setup process may be required prior interaction, which some users might experience as daunting or time-consuming. Also, certain connections, such as 3G, may incur costs, e.g., roaming fees. Moreover, network-based communication usually assigns a unique identifier to each communication participant, such as MAC addresses. Users and their devices can thus be tracked whenever they are connected to such a network, which may be regarded as a privacy threat.

Visual interaction addresses these challenges. The approach relies on a lightweight hardware and software infrastructure: A mobile application shows dynamically configured user interfaces and uses the built-in flashlight of a smartphone to emit sequences of light signals. Each sequence represents an action, e.g., selecting an item. A camera attached to the public display captures the light signals and translates them into application-specific actions, e.g., keystrokes. The set of available actions can be dynamically adapted. Visual interaction can thus be used as a generic tool to implement remote interactions with public displays via smartphones. The remainder of this subsection first presents related work on immediate and remote interaction

with public displays. Afterwards, the approach of visual interaction is discussed in more depth. Based on a prototypical implementation of the approach, as presented in Subsection 11.2.3, a study was carried out to assess the properties and limitations of the approach, see Subsection 12.2.3. The results provide evidence for the feasibility of optical interaction between smartphones and public displays.

Immediate Interaction

This type of interaction most frequently relies on physical contact or devices attached to the public display. This common type of interaction often uses prevalent and well-known devices such as keyboards or touchscreens. *Dynamo* [103], for example, is a public interactive surface, which multiple users can interact with in parallel via keyboards and mice. *Opinionizer* [188] enables people to place comments on a public screen via one public keyboard. The *CityWall* [180] can track an arbitrary number of hands and fingers on top of its surface by using high-resolution and high-speed cameras. The *UBI-hotspots* [130] can be controlled via touchscreens and can handle multiple users in parallel due to a spatio-temporal screen real estate management.

All these systems use physical input devices, which suffer from inherent problems: They have to be touched or pushed, for example, and are thus subject to misuse and vandalism [148]. Also, most people use their bare hands to control these devices. This raises hygiene concerns [73, 120], as there may be a risk of spreading germs. Privacy issues may also result from this interaction: Malevolent users may, for example, spy on others while they type in their password [68].

Gestures could help to alleviate these problems. Vogel and Balakrishnan [231], e.g., investigated hand gestures to control personal content on screen. Nancel et al. [160] compared different mid-air gestures in

front of wall-sized displays. They suggest to avoid gestures performed in free space as they tend to be less efficient and more fatiguing than gestures on surfaces. *WaveWindows* [182] allowed users to interact via waving or knocking. Study results indicate, however, that social inhibition may prevent users from performing such gestures.

Gesture-based approaches may thus share common advantages and drawbacks. In most cases, cameras track the gestures. These cameras are usually secured in protective casings. Thus, there is no exposed physical interface that could be vandalized. Moreover, gestures may be well-suited in scenarios in which hygiene is an issue. However, gestures are a quite novel approach. Many user may be unfamiliar with this means of interaction and have difficulties or even anxieties to use it, e.g., due to a lack of unified gestures [120] that may confuse users [234]. Finally, gestures may raise privacy concerns [182]. For example, due to current technical limitations, gestures have to be quite distinct and thus easily observable, so that they may be unsuited for sensitive data, such as passwords [43, 123]. This may also cause social inhibition.

Remote Interaction

Remote interaction is an alternative means of interaction with public displays. Though users do not have to be within arm's reach of the display, the interaction is often limited to a short range, e.g., a few meters. This constraint appears natural, as public displays are often situated at specific locations [174], and larger distances between the display and its users could negatively impact this characteristic. *Sweep and point and shoot* [23] lets users select objects on screen and drag them around via smartphones. Though the set of supported actions, e.g., selecting and activating objects, is comparable to the one presented in this subsection, it only provides one generic action set.

SnapAndGrab [137] facilitates bidirectional information transfer between public displays and smartphones. Users may send pictures of interesting screen areas to the display, which responds with a downloadable “data package.” Crossmodal Displays [167], e.g., the Rotating Compass [189], use smartphones to issue tactile cues when content is shown that could be relevant to a specific user. The two approaches presented by Boring et al. [33] and Baldauf et al. [21] let users employ their smartphones to project interactions issued on the display of the smartphone onto the public display. In contrast to visual interaction, all systems mentioned above require a network connection, e.g., 3G, WiFi, or Bluetooth, between the public display and the smartphone.

Though there may be some advantages to remote interaction in comparison to immediate interaction, there are also potential drawbacks: The awareness of privacy issues in the general public has increased in recent years, e.g., in the aftermath of disclosures on governmental surveillance. Privacy has in turn become a sensitive topic with respect to Internet-based communication. Interacting with public displays over the Internet may thus appear as a privacy threat to some users. Optical communication could be a means to overcome this issue. Screen Codes [58], for example, facilitate the transfer of data from public displays to hand-held devices. In contrast to this unidirectional approach, FlashLight [94] uses the smartphone camera to receive data and the built-in flashlight to send data to the public display. A key difference to visual interaction is that FlashLight requires the smartphone to be placed right on top of the display surface. C-blink [147] can be used to remotely control public displays via optical signals shown on the screen of a mobile phone. Yet, for some functions, the system requires an additional network link. Moreover, C-blink is limited to one generic action set, while visual interaction, as proposed in this thesis, can be tailored to different application scenarios.

Flashlight interaction [204] uses the flashlight of the smartphone instead of its screen: Users control a cursor by moving their smartphone in front of the public display. This has some resemblance to the approach presented in this subsection. The difference is, however, that Flashlight interaction only supports one action set, i.e., pointing, clicking, and zooming, whereas visual interaction supports various action sets that match individual application scenarios. *Lumitrack* [243] is an optical motion tracking system that requires dedicated hardware to project and detect special visual patterns, called *m-sequences*. *SideBySide* [240] allows for multi-user interaction with handheld projectors. The projected contents, e.g., animated characters, may interact based on information transferred via infrared light.

Overall, the review of previous work shows that none of the proposed remote interaction techniques are without issues. Compared to many of the conventional network-based communication approaches, systems based on optical communication may provide an increased protection of privacy, as the communication between public displays and smartphones is local, ephemeral, and may be harder to capture if not collocated. Most of the optical systems presented above do not require users to complete a coupling or pairing process, e.g., entering credentials or exchanging security tokens. Users may thus interact instantaneously with these systems. Pure optical communication also avoids additional costs, especially roaming fees.

Proposed Approach to Visual Interaction

The basic design of the approach consists of two components: a public display (CMP1) and a smartphone application (CMP2), see Figure 10.12. The smartphone application allows to remotely control the public display via a user interface that is tailored to suit a specific application

scenario, i.e., the application running on the public display. The appearance and behavior of the user interface depends on the specific public display and is dynamically transferred from the public display to the smartphone. The communication between both components is based on optical means alone as explained below. The public display is equipped with a camera, while the smartphone application requires a built-in camera and flashlight to receive and transmit optical information. In simplified terms, the smartphone is comparable to a TV remote control, except that the UI of the remote is adaptable, and no infrared (IR) light—and thus no special hardware—is used.

Figures 10.12a–d illustrate the basic design of the proposed approach and visualize the workflow, which was also implemented in the prototype, see Subsection 11.2.3. In the sketched example, users can perform one of three actions by pressing the corresponding button (here, represented as white, gray, and black) on their smartphone. However, the user interface is not limited to this particular UI, it merely serves as an illustrative example. As depicted in Figure 10.12a, the smartphone application dynamically retrieves data about the UI from the public display via a QR code.

QR codes were chosen, since they can be easily integrated in existing public display software (e.g., as JPEG images) and can be scanned reliably from various perspectives, see Subsection 10.2.1. Moreover, many people are familiar with QR codes, as they are a common sight on billboards, posters, etc., which could help to reduce the initial interaction barrier from the user’s point of view. Since QR codes can contain any type of data, they can also transfer data about a user interface that suits a specific public display application. One way to convey this data would be to use the XML User Interface Language (XUL⁷). Figure 10.12b illustrates how the smartphone application extracts the

⁷<https://developer.mozilla.org/de/docs/XUL>, accessed: July 15, 2015

UI as well as the corresponding coding scheme from the QR code. Later, the coding scheme is used to encode light messages emitted by the smartphone flashlight.

Figure 10.12e shows three examples for user interfaces suitable for different scenarios: The colored buttons (top) let users select items that they would like to get additional information about (for example, pressing the white button would show world news, the gray button sport news, and the black button weather forecasts). The d-pad (middle) could be used to let users move a game character. The vertical slider (bottom) could be used to control the zoom level of a map. As mentioned above, the design of the approach also lets users directly control an onscreen cursor or pointer, which is comparable to the sweep method introduced by Ballagas et al. [23]. However, this action set was not included in the prototype, since Ballagas et al. already conducted a similar study. When the user performs an action, e.g., pressing the black button, the smartphone application optically transfers this information to the public display by emitting a light signal, see Figure 10.12c. Eventually, the public display shows the selected content, see Figure 10.12d.

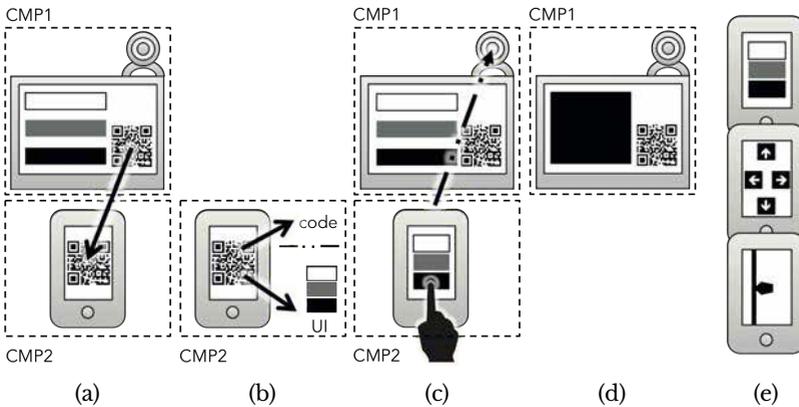


Figure 10.12.: The design of the visual interaction approach with its components CMP1 and CMP2. (a) First, the smartphone application (CMP2) photographs a QR code shown on the public display (CMP1); (b) the QR code contains data about the coding scheme and the UI; (c) the transmitted UI is shown on the smartphone and the user interacts with it. User actions are encoded using the transmitted scheme and sent to the public display via the flashlight of the smartphone. The camera of the public display then captures the light signals; (d) the public display decodes the signal to identify the user action and triggers the corresponding application-specific action; (e) three examples for adaptive user interfaces, i.e., color code buttons, d-pad, vertical slider (top to bottom).

10.3. Process Integration

As presented in Chapter 9, there are at least eight challenges that make the design of public display systems a complex task. In particular, public displays often allow for interaction, for example, via touch-screens, WiFi and Bluetooth interfaces, or body tracking cameras. These means of interaction add complexity to the software that runs the public display and to the user interfaces. As with any software, public display systems need sound design and testing to guarantee a certain level of reliability and to ensure a good user experience. There are best practices and approaches, e.g., unit testing and participatory design, to develop and test conventional software. Yet, many of these practices only partially address key aspects of public displays, i.e., their situatedness and their inherent interaction with the context they are deployed in.

Alt et al. [12] compared various evaluation methods for public display systems and analyzed them with respect to how well they are suited to answer specific research questions, for example, about the audience's behavior or the user acceptance. One aspect they highlighted is the importance of ecological validity for public displays, which usually comes at the expense of external or internal validity. This section thus proposes a novel approach to increase the degree of situatedness in lab-based public display studies.

The approach supports the design and evaluation of public display systems at early development stages by combining panoramic images or video footage with a light-weight, graph-based model to simulate public displays. Subsection 11.3.1 presents the *Immersive Public Display Evaluation and Design (IPED) Toolkit* as a prototypical implementation of the approach. An *Immersive Video Environment (IVE)*, see Subsection 11.3.2, allows users to experience the public display system as

if they were exposed to a real, physical deployment. Benefits of the toolkit in terms of rapidly designing and evaluating public display systems with relatively little effort are shown.

10.3.1. Immersive Public Display Evaluation and Design Toolkit

First, this subsection positions the proposed approach within related work. In doing so, it addresses the design, prototyping, and evaluation of public display systems. Afterwards, the actual approach towards a process integration (C3) is proposed as the IPED Toolkit. The corresponding prototype is presented in Section 11.3. Finally, Section 12.3 reports on experiences gathered from using the proposed approach in a real application scenario; moreover, the same section reports on results obtained from a user study, assessing the suitability of the proposed approach to address each of the eight challenges.

Design of Public Display Systems

A fundamental property of signage in general is its legibility. Xie et al. [244] investigated (static) emergency signage and its legibility. They proposed and validated a geometrical model that captures relevant aspects of sign visibility. Such an approach could be applied during the design phase of a public display to assess and predict whether the shown content can be perceived. The same is true for agent-based models [181], which can simulate pedestrian movement in general, but can incorporate visibility assessment as well. As such models often only rely on geometrical computations, combining them with the approach presented in this thesis would facilitate the investigation of contextual factors as well.

Stahl and Hauptert [214] used 3D models of intended deployment sites in which they inject screen contents of existing public display systems. This is similar to the approach proposed here. However, their system lacks visual fidelity, does not incorporate sensor readings, and does not easily support different scenes at the same location, for example, daytime vs. nighttime. A (scaled-down) replica of a planned deployment site can also be used in the design phase. Hamhoun and Kray [90] applied this approach to a public display system that supports navigation at densely crowded sites. They were able to gather insights into properties of the full-scale system, e.g., the relative density of displays, based on the physical simulation. While the approach proposed in this thesis could be used to design and evaluate the system presented by Hamhoun and Kray, simulating the users' locomotion is still difficult.

Prototyping Public Display Systems

The Proximity Toolkit [133] is based on a theory of proxemic interaction. It allows for rapid prototyping of applications that use the ideas and concepts of proxemics, e.g., the user's location and orientation in front of public displays. While this approach was successfully used to prototype different public displays, it only focuses on one type of interaction and does not consider contextual factors, display contents, or the deployment location of the public display.

This last aspect is picked up on by Nakanishi [159], who proposed to use miniature models of real locations to facilitate frequent prototyping and testing. Not all relevant aspects can be covered in a miniature model, which is why he suggested to analyze a corresponding virtual model as well. This virtual model can then be used to assess the ideal positioning of interactive devices, while the miniature model can be

used to eliminate discrepancies between the virtual and the real space, such as optical attenuation. His approach is subject to some limitations, as the virtual model and the small scale physical model may lack some realism. Moreover, in contrast to the approach presented in this thesis, his system requires designers to work with two models rather than just one.

Harrison and Massink [92] proposed stochastic models as a means of prototyping and evaluating ubiquitous systems prior to deployment. The idea is to reduce risks during the development, as some design flaws can be identified early. However, constructing usable models of this type may require considerable effort and might not capture all relevant contextual factors. In contrast, the approach proposed in this thesis facilitates a rapid development of simulations and integrates different contextual factors.

The APEX framework [206] is a related approach for model-based rapid prototyping of ubiquitous environments. It enables users to experience an envisioned system in a 3D simulation and is based on three components: a virtual environment, a behavior, and a communication/execution. These components bear some resemblance to the proposed architecture described below. However, the model of the framework is based on a—potentially complex—Coloured Petri Net (CPN) in contrast to the light-weight, graph-based model presented in this thesis.

Evaluation of Public Display Systems

Alt et al. [12] presented a survey on how to evaluate public display systems based on an extensive literature review, see Section 6.4. They identified a set of typical research questions that often occur when evaluating public displays and classified how such research questions

could be evaluated. They also discussed external, internal, and ecological validity and provided a small number of guidelines for studying public display systems with users. Though they did not consider Immersive Video Environments or formal models, they covered a broad range of evaluation methods and highlighted the relative benefits and drawbacks of different methods.

One of the most popular evaluation methods they identified was to record all interactions, for example, via log files or video recordings. This approach has been used in many different settings [65, 79, 217, 225], in particular in combination with extended deployments such as reported for the WrayDisplay [223, 224], the *Hermes* system [51, 52, 75, 117, 220], SPAM [50], and MobiDiC [153]. A key advantage of an approach based on recordings is a high degree of ecological validity, since the interactive public display system is analyzed in its target environment. Drawbacks include privacy concerns, inherent limitations in terms of what can be recorded, as well as the effort and time required for long-term deployments. The approach presented below can make use of such recordings to simulate physical environments in the lab and also supports recording interactions there.

Singh et al. [208] combined immersive video and surround audio to create “a realistic simulation of a ubiquitous environment” [208]. The toolkit presented below was inspired by the design of Singh et al. However, while Singh et al. focused on prototyping and evaluating context-aware apps on mobile devices, the toolkit proposed here supports the design, prototyping, and evaluation of public display systems. All of the presented methods for design, prototyping, and evaluation of public displays have their strengths and weaknesses. Table 10.15 summarizes the characteristics of each method and contrasts them with the approach described in this thesis.

Table 10.15.: Overview of evaluation and design methods for public display systems.

Author (system)	Pros	Cons
Marquardt (Proximity Toolkit) [133]	Rapid prototyping possible	Limited interaction; lack of context
Hamhoun & Kray [90]	Supports locomotion	Effort required to construct physical small-scale model
Harrison & Massink (PEPA, Fluid Flow) [92]	Prototyping and a priori evaluation based on stochastic model	Effort required to create model; may require expert knowledge; lack of context
Nakanishi [159]	Rapid prototyping of interactive public display systems	Requires analysis of miniature and virtual model rather than just one; interaction limited
Taylor (WrayDisplay) [223, 224], Faisal (Hermes) [220, 117, 75, 51, 52], Cheverst (SPAM) [50], Müller (MobiDiC) [153]	Results with high external and internal validity	Inherent limitations of logging; effort and time required for long-term deployments
Stahl & Hauptert [214]	Rapid evaluation of display and content visibility	Limited coverage of other design and evaluation aspects

Table 10.15.: Overview of evaluation and design methods for public display systems (continued).

Author (system)	Pros	Cons
Silva et al. (APEX) [206]	Precise (3D) model of evaluated system	Based on (complex) Coloured Petri Net (CPN)
Singh et al. [208]	Rapid and low effort prototyping of context-aware mobile apps; no physical deployment needed	Applicable to mobile apps only; not accessible to non-experts; locomotion limited; predefined locations; limited applicability to research questions defined by Alt et al. [12]
IPED Toolkit	Rapid and low effort prototyping of interactive public display systems; accessible for non-experts; no physical deployment needed	Locomotion limited; predefined locations; limited applicability to research questions defined by Alt et al. [12]

Core Elements

The approach to design, prototype, and evaluate public display systems as presented here aims at replicating real-world scenes in the lab. A high degree of audiovisual fidelity can provide designers and participants with an immersive experience similar to being in situ. The approach consists of a number of core elements.

The central element is a state-transition graph that encapsulates the different states which the simulated world can be in, e.g., locations or daytime vs. nighttime. This graph-based approach was chosen because of its light-weight and extensible characteristics. Moreover, traversing the nodes in the graph mirrors the physical transition between locations and situations in the real world. Locations are usually situated in a specific area of interest, i.e., the area where a public display system is meant to be deployed. It seems reasonable to focus on these locations, as the inclusion of all possible locations would result in a very large—and cumbersome—graph.

Instead, the approach focuses on decision points and places, which have a specific relevance in the investigated application scenario (third challenge, Section 9.3). Consider, for example, a public display that facilitates the use of a public transport network in a city. It would be reasonable to include locations that are served by public transport in the graph, for example. A single physical location can be represented by more than one node as explained below, e.g., to capture different states or contexts of this location.

The edges in the graph represent transitions from one node to another. As each node corresponds to a location or state, an edge basically determines whether it is possible to directly move from one location or state to another, since one location can be represented

by multiple nodes to encode different states (fourth challenge, Section 9.4). In the public transportation scenario, an edge might connect two nodes that represent two adjacent bus stops on a particular bus route. Equally, an edge might connect nodes representing different states of the same location, e.g., one node represents the location while a bus arrives and another node represents the same location with no buses at all.

The graph describes physical as well as logical connections, events, and other links between different scenes. Designers can use this structure to describe specific use cases or scenarios for their public display system. Users can experience the virtual world by moving through the graph. Users can only be in one location or state at a time, i.e., they can only be at one node of the graph. The graph thus represents the envisioned installation sites of public displays by connecting users, displays, locations, and states.

The basic graph structure also provides a framework to organize and attach audiovisual and other additional data to locations. Moreover, it allows for the integration of public displays in the simulated world. Audiovisual material, such as video footage, photographs, or audio recordings, can be captured at the locations of interest and can then be used to simulate locations and their states during design, prototyping, and testing. In the context of the public transportation scenario, for example, several short video clips could be recorded at a bus stop for each of the relevant states, e.g., one showing a bus arriving, one showing the stop without a bus in sight, and one showing the bus departing from the stop.

Further data can also be recorded on site, such as GPS and orientation information or environmental factors, such as temperature or signal strengths of cell towers. This data can be linked to a specific node in the graph and can then be used during the development process.

In the context of the bus stop scenario, the GPS and orientation data could be used to design and evaluate the behavior of an augmented reality application that visualizes actual bus routes and departure points on the screen of the user's smartphone, taking into account which information is being shown on a public display at the bus stop.

Public displays are specified within the frame of reference defined by the visual data linked to a location. For example, regions in the footage can be labeled as public displays. During the design, prototyping, and evaluation, they can be replaced by the envisioned content. The proposed approach does not specify how the content is generated, but only how it is integrated in the simulated world. Content can be generated in a variety of ways, e.g., Wizard-of-Oz style (WoZ), static imagery, a functional system that serves content adapted to various locations, etc. The public displays can be placed via the frame of reference within the visual data in different ways as well. A simple option is to use the pixel-coordinates of a video frame or photograph. Alternatively, a set of depth layers can be defined on top of a video scene or photograph. Public displays could then be placed on a specific layer. Finally, a 3D model can be associated with the visual material. This 3D model describes all visible surfaces, e.g., walls or tables, geometrically; public displays can then be attached to these surfaces.

For example, if the video footage shows the view from a bus stop looking at buildings across the street, a simple 3D model could include the facades of those buildings. Public displays could then be attached to one of the virtual facades at a particular location (first and second challenge, Sections 9.1 and 9.3). The frame of reference is also important to facilitate the interaction with simulated public displays: Within a specific scene, the relative location of a person in front of a (virtual) public display can also be specified within the given frame of reference. It is thus possible to realize distance- or orientation-based in-

teraction with the virtual public display, see Subsection 11.3.2. For example, the content of a time-table display at a bus stop may change depending on how close a person stands to the display.

Creating Simulated Environments

Unlike fully synthetic simulations, e.g., 3D renderings based on textured geometric models, the proposed approach requires considerably less effort to generate realistic simulations while providing means to easily modify and test key characteristics of a public display system. The following paragraphs describe the process of creating such simulated worlds in more detail.

The creation of simulated worlds consists of five steps, see the boxes in the middle row in Figure 10.13. If the simulated world should include sensor data, e.g., GPS or compass information, three additional steps have to be performed, see the boxes in the top row of Figure 10.13. The first step is to identify decision points that have a specific relevance in the analyzed application scenario. Looking at the public transportation scenario, for example, the relevant decision points could be the stops of a particular bus line. The next step is to construct the graph with its nodes and edges. A node represents a particular decision point in a specific state. An edge represents a possible transition between two decision points or a transition between two states of one particular decision point.

The third step is to record the actual decision point in situ. Depending on the hardware being used and the goals being pursued, the recording can be done with one or more (video) cameras, audio capturing devices, or other sensors. Once the recording is complete, the footage has to be post-processed. This may include steps such as adjusting the resolution, performing panoramic stitching, format conversions, or

creating seamless video loops. The final step is to link the recorded footage to the corresponding nodes in the graph.

In case the recorded footage is complemented by sensor data, the recording step has to be extended with appropriate devices, e.g., a GPS tracker or a compass. This additional sensor data may also require post-processing, for example, to align the measured samples to certain time codes in the video footage. Finally, the sensor data needs to be linked to the corresponding nodes in the graph as well.

Integrating Public Display Systems

In order to integrate public display systems in the simulated world, designers have to carry out five steps that correspond to the five steps explained above. The steps are depicted in the boxes in the bottom row in Figure 10.13. The first step is to specify the placement of a public display at designated decision points. This is important since the placement may have some influence on the actual video recording, e.g., in terms of distance or perspective. The next step defines the screen content and ways of interaction with the public display. In terms of the local transportation scenario, this could be to define whether the simulated public displays show the bus schedule of a specific route or rather instructions on how to interact with this schedule.

The third step in the process is to create the screen content. This can be done in different ways, for example, by using fully functional systems, prototypes of varying fidelity, or simple (static) mock-ups. The fourth step is to define the display overlays. This includes, for example, to specify the exact position and spatial dimensions within the frame of reference. Finally, the fifth step links the overlays to the nodes in the graph so that the simulated public display system is shown whenever the user arrives at the corresponding node. In

practice, most of the steps may take little time. Post-processing the footage, for example, could only require to crop it so that it can be looped infinitely.

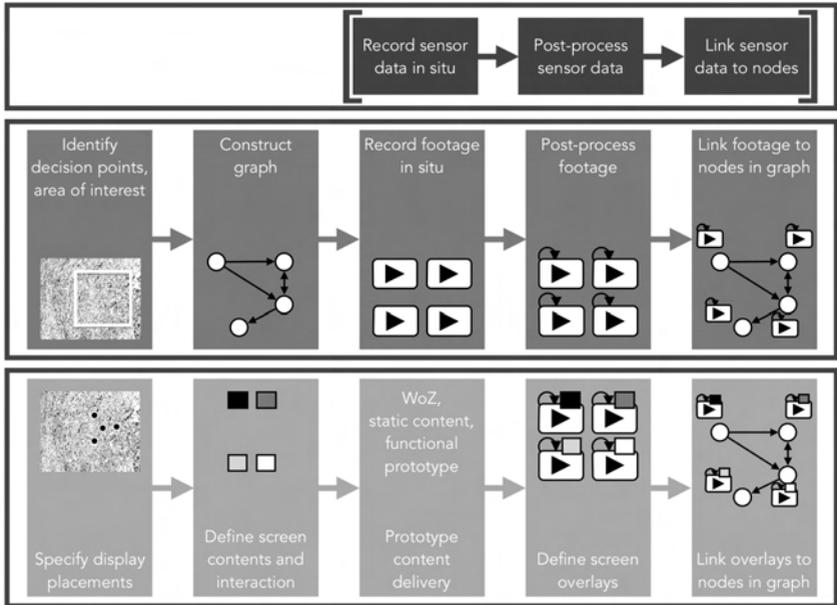


Figure 10.13.: Process of preparing the simulation environment and integrating public display prototypes. The first two rows illustrate the construction of simulated environments; the bottom row outlines how to integrate a public display into such environments.

Example Application

The proposed approach can be applied to different phases of the development. The following paragraphs review how it can be used during the design, prototyping, and evaluation of public display systems.

Design. The design of an interactive public display system is a complex task that is influenced by many aspects such as, e.g., the location, orientation, form factor, background, or content of the public display, cf. the eight challenges introduced in Chapter 9. The approach proposed in this thesis facilitates the manipulation of multiple parameters with ease and at low cost. In terms of the local transportation example, it would be possible to adjust the height of a public display effortlessly until it suits the needs of, e.g., physically impaired people. Also, the visual appearance of a public display could be easily altered in order to determine the size, color, or shape that attracts people the most. The approach thus allows for user-centered design (UCD) or participatory design, for example, by engaging multiple stakeholders in discussions while experiencing the simulated environment. Similarly, sensitive content, e.g., personalized information, can be analyzed and immediately revised in a realistic simulation of a public environment. This can also support the legal assessment prior to public exposure. Thus, the proposed approach can help to address the first four, the sixth, and the eighth challenge.

Prototyping. Closely linked to the design process is the prototyping of a public display system. The proposed approach supports this phase with realistic simulations of the behavior of interactive public displays. Looking at the bus stop example, the approach could help to prototype a multi-display network that is spread throughout the

city to deliver up-to-date arrival and departure times. The graph of the simulated world could be used to let users take a virtual walk—or bus ride—through the city in order to test the interaction between the users and the interactive public displays, e.g., via touch screens or different mobile devices. Thus, the proposed approach helps to address the fifth and sixth challenge.

Evaluation. The graph, which is at the heart of the proposed approach allows for manipulating various parameters such as the location of the public display, its orientation, content, or interaction mode. The approach can thus complement conventional controlled lab tests or field studies. Previous research shows that people who are exposed to an IVE can actually feel immersed under certain conditions [210]. This facilitates the evaluation of non-functional characteristics of an interactive public display system. Regarding the public transportation scenario, the proposed approach could thus be used to evaluate how users perceive a multi-display network depending on its presence at certain (sensitive) points of interest within the city.

10.3.2. Immersive Video Environment

Video- or photo-based environments have been used in the past to evaluate situated technology, in particular mobile systems. Snowdon and Kray [210], for example, used such an environment to assess a mobile system that provides hikers with information about natural environments. Even without sound or moving images—the simulation used panoramic photographs only—they reported on a high degree of immersion amongst the participants, evidenced by the way in which they referred to objects depicted on the screens. In particular, the language people used to describe the scene and to identify and

locate objects shown on screen strongly resembled what they would use if they actually were at the location in the real world.

For example, in some scenes participants would refer to objects as if standing at the top of a hill looking down into a valley rather than just describing what they see on a screen. It therefore stands to reason that certain contextual factors, such as the structure of the environment where a public display is installed, can be replicated well inside an Immersive Video Environment. Thus, it can also be used during the development and evaluation of public displays when applying the approach proposed in this thesis.

There are different approaches to realize visually convincing simulations. One is to use virtual environments (VE), which are computer-generated scenarios based on elaborated 3D models of given situations. Another approach is to use photographs or video footage to generate an immersive experience. Each approach has different advantages and drawbacks. Synthetic 3D models allow for fine grained details and interaction, while the actual modeling requires a lot of work. While parts of the process can be automated, e.g., using 3D scanners [104], the overall effort is still considerable.

Conversely, photographs or video footage provide realistic (audio-) visual simulations and can be captured quite easily. Yet, interaction with them is limited compared to a 3D model, in particular with respect to locomotion, but sufficient in the context of this thesis. Lee et al. [127] propose a virtual reality environment to systematically compare augmented reality applications. They show that such a simulation may provide a reasonable validity, as long as the latency of the system is kept at a minimum. The idea presented by Lee et al. is comparable to the one discussed in this thesis. The implementation, however, differs as Lee et al. use computer generated graphics rather than video footage.

Locomotion is another factor contributing strongly to creating convincing simulations. Different means exist to realize this, e.g., treadmills, gloves, artifacts, or gestures. Uni-, bi-, or omni-directional treadmills can be used to let users actually walk in a virtual environment [55]. Schellenbach et al. [197] discovered that users' walking patterns on treadmills match those of walking overground quite well, provided that users are given enough time to get accustomed to the system.

Gloves allow for direct interaction and manipulation of virtual objects and can also provide haptic feedback [36]. Kim et al. [113] propose a system that does not require users to wear gloves, but wrist-worn sensors. These sensors optically analyze the pose of the user's hand and process the result into kinematics models. While these approach can only mimic the original sensation, physical artifacts are another means for manipulating a virtual environment that provide users with an authentic haptic sensation [177]. Pushing physical buttons or turning knobs, e.g., will trigger certain actions in the virtual environment.

Gestures can also be used to manipulate virtual objects. Due to recent advances in consumer electronics, such as the Kinect camera by Microsoft, research on gestures has attracted a lot of attention. Gestures have the potential to provide intuitive access to a virtual environment, and have consequently been a subject of investigation in this area [25, 160, 193, 232]. However, since humans perceive visual stimuli more pronounced than auditory or tactile ones [14], the approach presented in this thesis focusses on the visual simulation within an Immersive Video Environment.

As a survey of publications on mobile HCI during the last decade revealed, there has been a shift from engineering-driven to empirical, evaluation-based research [115]. An Immersive Video Environment provides a way to control the context of use and thus can be used to facilitate empirical research, e.g., for systems supporting mobile navi-

gation [208, 210]. An Immersive Video Environment can also provide a useful platform to investigate aspects pertaining to other types of context-aware systems, such as display blindness or interaction blindness [156, 164], with regard to public displays.

A common physical setup to realize an Immersive Video Environment is a CAVE. User studies demonstrated that users actually feel immersed, i.e., they experience the situation as if it was real [210]. Given these properties, it makes sense to also consider the use of Immersive Video Environments to design, prototype, and evaluate public display systems. Kray and Delikostidis evaluated location-based services (LBS) in-the-field as well as in-the-lab, based on an Immersive Video Environment, and compared both approaches [71, 70, 118].

While Immersive Video Environments are easy to create and very realistic, they do not include semantic or geometric information. Movement in the depicted 3D space and interaction with objects shown in the footage is thus not realized easily. The approach presented in this thesis combines gestural interaction with a mirror image of the user that serves as an avatar within the video environment. It thereby enables the intuitive selection of 3D locations shown in video environments as well as the placement of virtual objects inside the 3D space depicted by the video footage.

Similar to the approach presented by Ahn et al. [4], the approach proposed in this thesis uses the user's mirror image as a video avatar. The avatar can be used to navigate within the virtual environment and to manipulate virtual objects. In contrast to previous work, this approach does not merely substitute a cursor with a mirror image. The approach is focussed on providing an immersive experience by using gestures in combination with the video avatar. It supports both the creation of augmented scenes, where virtual objects are inserted into video footage and the exploration of such scenes.

Gotardo and Price [85] aimed at a similar workflow and developed a system that is comparable to the one presented here. However, their approach is based on a more complex hardware setup. Moreover, the user interface designed by Gotardo and Price is based on heads-up displays rather than on gestures to let users select or scale objects, for example. A purely gestural interface may be experienced as a more natural way to interact with an Immersive Video Environment, both by designers and participants. In addition, the system proposed here allows for a quick and easy creation of scenes from video footage and does not require custom and expensive hardware. Subsection 11.3.2 presents more details on the prototypical implementation.

11

Prototypes

As explained in Section 4.2, this thesis strives to provide tangible scientific contributions. This chapter thus presents prototypical implementations of the privacy threat model (C1), see Section 11.1, the three novel countermeasures (C2), see Section 11.2, and the process integration (C3), see Section 11.3. The underlying intention is twofold: Firstly, this allows for a well-founded evaluation of the particular approach, for example, based on user or field studies, see Chapter 12. Secondly, other researchers or designers of public display systems may directly use the scientific contributions C1–C3, possibly serving as a springboard for future work.

11.1. Privacy Threat Model

Though the privacy threat model proposed in Section 10.1 is useable per se, researchers as well as designers may be uncertain about how to actually apply it. There might be the desire for a guided process, that helps to govern the analysis of privacy threats in a systematic and structured way. In order to address this desire, facilitate the evaluation of the privacy threat model (C1), and to provide the model in a tangible form, the findings were incorporated in a web-based tool,

called the Interactive Public Display Privacy Threat Model (*IPDPTM*). Figure 11.1 shows a screenshot of the publicly available¹ tool. As explained above, the prototypical manifestation of the privacy threat model also allowed for a well-founded evaluation, see Section 12.1.

The tool can be used by researchers, developers, and designers of interactive public displays to design privacy-preserving systems. Based on technologies such as HTML5 and D3², it allows users to save, load, and export threat models. Exported threat models can be re-used as vector graphics in various types of software and documentations, e.g., functional specifications. This way, the IPDPTM can be integrated into existing processes, see also Sections 10.3 and 11.3. Furthermore, the use of established technologies, such as D3, allows the tool to flexibly adapt its behavior and appearance in specific application scenarios. Figure 10.4, for example, was also created with the IPDPTM.

The web-based tool helps to structure all related entities by visualizing the results and providing textual examples. In doing so, the tool guides the modeling process by showing all applicable options in a certain situation. For example, all reasonable (i) threat types (depicted as red squares in Figure 11.1) with regard to specific threat agents (depicted as purple squares in Figure 11.1) and all potential (ii) weaknesses (depicted as yellow squares in Figure 11.1) according to that particular threat type, as well as all possible (iii) effects (depicted as blue squares in Figure 11.1) along with (iv) suitable countermeasures (depicted as green squares in Figure 11.1), which may be used to mitigate that weakness. The countermeasures proposed by the tool are based on the review of related work, see Subsection 7.5.3.

In order to show only the relevant or applicable options, countermeasures to a specific threat and weakness, for example, the IPDPTM ana-

¹<http://ipdptm.se-labor.de>, accessed: May 21, 2015.

²<http://d3js.org>, accessed: May 21, 2015.

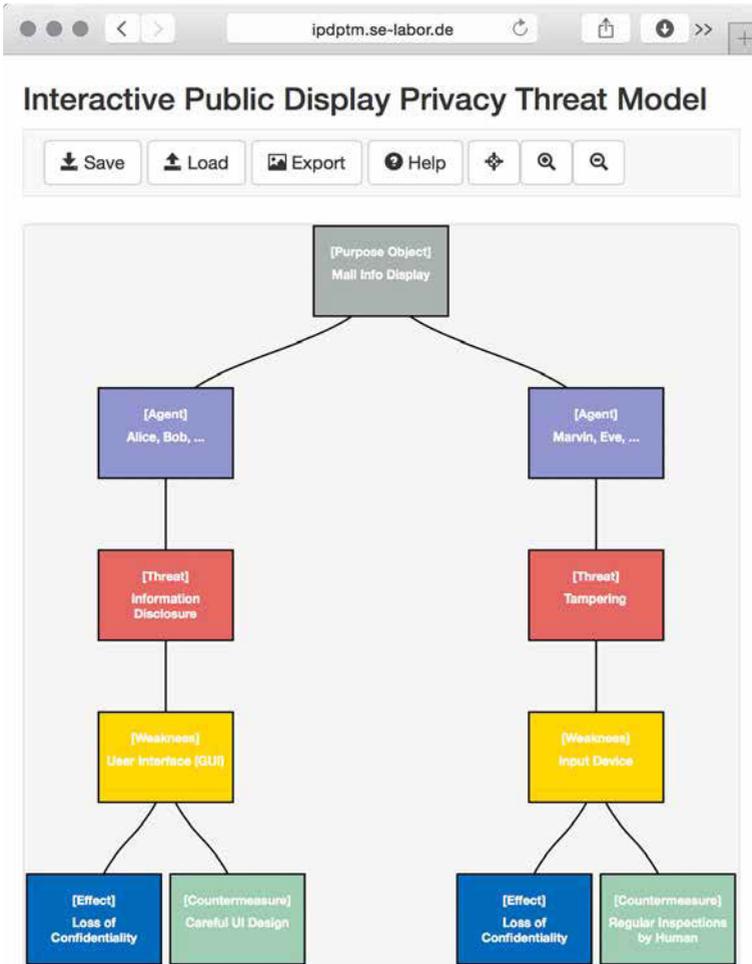


Figure 11.1.: Screenshot of the Interactive Public Display Privacy Threat Model (IPDPTM).

lyzes the relations defined by the privacy threat model. The complete list of relations can be found in the appendix on pp. 444. Each bullet point shown there directly correlates to one item in a JSON array that is parsed by the IPDPTM. To account for situations in which the pre-defined relations may not apply, users may add individual items, such as threats, for example.

The following example illustrates the application of the prototype: At first, the web-based tool presents a default purpose object, named “public display.” Users can now click on that purpose object to open a context dialog. This dialog allows them to specify the name of the node, e.g., “Mall Info Display” (cf. Figure 11.1), or to add an arbitrary number of child nodes, i.e., threat agents. After adding the child nodes, users may click on these child nodes in turn to open their context dialogues. Here, users can select the name of the threat agent from a list of pre-defined options, e.g., “Alice, Bob, ...,” or they can specify a custom name. To help users find the desired list item, users can hover their mouse over a particular item; a tooltip then provides some examples to guide users in their decision. Furthermore, users can also add an arbitrary number of child nodes, i.e., threats.

After adding the threats, users can open the corresponding context dialogues. Again, users can choose the name of the node from a pre-defined list. This time, however, the available list items depend on the name of the preceding threat agent, since not all agents can perform each threat. Users can also add an arbitrary number of child nodes, i.e., weaknesses. Users can now use the same process to name weaknesses and to add child nodes. Weaknesses, however, can have two types of child nodes, i.e., effects and countermeasures. The former ones can not have any child nodes in turn, but users can add another purpose object as a child node to each countermeasure. This allows for realizing the cyclic iterations described in Subsection 10.1.2.

Combining the tools and processes presented in Sections 10.3 and 11.3 with the IPDPTM can result in a holistic tool for privacy-preserving public display systems. That tool can be used by researchers and designers to design and evaluate such systems.

11.2. Countermeasures

To contribute to the list of countermeasures identified in the review of related work, Section 10.2 proposed three novel countermeasures, i.e., visual multiplexing, visual highlighting, and visual interaction. This section presents the prototypical implementations of each countermeasure and discusses their individual characteristics.

11.2.1. Visual Multiplexing

As mentioned in Subsection 10.2.1, the concept of visual multiplexing is based on two components, i.e., a software that multiplexes the input images and a mobile application that demultiplexes the corresponding information channels. Accordingly, there are two prototypical implementations—both called *Multipleye*—, one for each component. The prototype for the multiplexer has been realized as a public web application³. The prototype for the mobile application has been realized as an iOS application, that runs on, e.g., iPhones or iPads from Apple. The application is also publicly available⁴. The three following subsections focus on the three presented multiplexing methods, i.e., frequency-division multiplexing (FDM), code-division multiplexing (CDM), and time-division multiplexing (TDM).

³<http://www.multipleye.de>, accessed: May 18, 2015.

⁴<https://itunes.apple.com/td/app/multipleye/id441059663>, accessed: May 18, 2015.

Listing 11.1: FDM multiplexer prototype: Algorithm used for additive color mixing in PHP using the GD library. The multiplexed image is named \$imageRGB.

```
for ($w=0; $w < $width; $w++) {  
    for ($h=0; $h < $height; $h++) {  
        $R = (imagecolorat($imageRed, $w, $h) >> 16) & 0xFF;  
        $G = (imagecolorat($imageGreen, $w, $h) >> 8) & 0xFF;  
        $B = imagecolorat($imageBlue, $w, $h) & 0xFF;  
        $rgb = imagecolorallocate($imageRGB, $R, $G, $B);  
        imagesetpixel($imageRGB, $w, $h, $rgb);  
    }  
}
```

Frequency-Division Multiplexing (FDM)

The web application allows to specify up to three individual texts. Optionally, users may rotate each text between -180° and 180° . Each text is used to render one input image, i.e., a red, a green, and a blue one. Eventually, the three input images are combined into the final multiplexed image by using additive color mixing. The color mixing is implemented with PHP and the GD library⁵, using the algorithm shown in Listing 11.1. Figure 11.2 is a screenshot of the web application.

The mobile application uses OpenGL ES 2.0 to manipulate the live camera video feed via a fragment shader. Users may select their desired information channels by pressing one of the three corresponding buttons at the bottom, see Figure 11.5. If the user selects the red information channel (Figure 11.5a), for example, the fragment shader turns on all pixels whose colors exceed a certain red threshold. All pixels that fall below that threshold are turned off. Users may adjust the threshold in the user interface (not depicted here); users may also select whether the activated pixels should be tinted in the color of the

⁵<http://php.net/manual/en/refs.utilspec.image.php>, accessed: May 18, 2015.

corresponding information channel or simply be black (also not depicted here). The corresponding algorithm is shown in Listing 11.2.

To adjust the mobile application to varying lighting conditions, e.g., outdoor vs. indoor use, users may calibrate the camera of the smart-phone. This calibration causes the camera to re-focus once and set the focus to fixed afterwards. This prevents unwanted losses of focus. Additionally, the calibration deactivates the built-in white balance in order to avoid an unwanted preprocessing of the camera image.

Code-Division Multiplexing (CDM)

The web application, see Figure 11.3, allows to define four input images. Users may upload arbitrary images, which will be resized and cropped to 128 x 128 pixels. The first version of the prototype (as used in the evaluation, see Subsection 12.2.1) then performs the steps explained on pp. 213. It uses OpenCV⁶ and the Fast Fourier Transform (FFT) algorithm provided by the C library FFTW⁷. Listing 11.3 shows the algorithm written as pseudocode for the sake of brevity. Once the strings for all input images have been created and concatenated to one string, that string is used to create the QR code via a web service⁸.

The second version of the prototype (not used in the user study) pursues a different approach. Rather than using an individual compression based on FFT, it employs the standard JPEG algorithm. As with the first version, the applied compression level is dynamically computed for every input image to allow for the highest visual fidelity while not exceeding the maximum QR code capacity.

⁶<http://www.opencv.org>, accessed: May 18, 2015.

⁷<http://www.fftw.org>, accessed: May 18, 2015.

⁸<http://goqr.me/api/doc>, accessed: May 18, 2015.

Listing 11.2: FDM demultiplexer prototype: Algorithm used to demultiplex individual information channels implemented as an OpenGL ES 2.0 fragment shader.

```
if (channel == 1) {
    if (pixelColor.r > threshold) {
        if (tintCameraImage == true) {
            gl_FragColor = vec4(pixelColor.r, 0, 0, pixelColor.a);
        } else {
            gl_FragColor = vec4(0, 0, 0, 0);
        }
    }
}

if (channel == 2) {
    if (pixelColor.g > threshold) {
        if (tintCameraImage == true) {
            gl_FragColor = vec4(0, pixelColor.g, 0, pixelColor.a);
        } else {
            gl_FragColor = vec4(0, 0, 0, 0);
        }
    }
}

if (channel == 3) {
    if (pixelColor.b > threshold) {
        if (tintCameraImage == true) {
            gl_FragColor = vec4(0, 0, pixelColor.b, pixelColor.a);
        } else {
            gl_FragColor = vec4(0, 0, 0, 0);
        }
    }
}
```

The mobile application uses the ZBar bar code reader library⁹ to read the QR codes generated by the web application. The QSUtilities¹⁰ are used for the base64 decoding. Users may switch between the available information channels by pressing the corresponding buttons at the bottom, see Figure 11.6. If the user selects the first information channel (Figure 11.6a), for example, the mobile application extracts the corresponding substring from the string contained in the scanned QR code. The first version of the prototype then uses the algorithm shown as pseudocode in Listing 11.4 to reconstruct the image. The second version of the prototype uses the JPEG algorithms provided by iOS to reconstruct the image.

Time-Division Multiplexing (TDM)

The web application, see Figure 11.4, allows to define up to four input images. Users may upload arbitrary images, which will be resized and cropped to 500 x 500 pixels. Though the size of each input image is not limited in theory, the prototype resizes and crops each input image to generate multiplexed videos with a homogenous look. The web application then uses FFmpeg¹¹ and ffmpeg2theora¹² to create the multiplexed video. For the prototype, the speed is set to 15 frames per second (fps), as the used iOS devices were equipped with cameras that allow for a maximum speed of 30 fps. According to the Nyquist–Shannon sampling theorem [229], this technical constraint implies that the multiplexed video may run at a maximum framerate of 15 fps. Despite these technical constraints, the current TDM prototype already allows to watch four TV stations in parallel on one

⁹<http://zbar.sourceforge.net>, accessed: May 18, 2015.

¹⁰<https://github.com/mikeho/QSUtilities>, accessed: May 18, 2015.

¹¹<https://www.ffmpeg.org>, accessed: May 19, 2015.

¹²<http://v2v.cc/~j/ffmpeg2theora>, accessed: May 19, 2015.

Listing 11.3: CDM multiplexer prototype: Algorithm used to compute the FFT data as pseudocode.

```

for every row and column in image_data {
  data[column][REAL] = image_data[row * width + column];
  data[column][IMAGINARY] = 0.0;

  computeFFT;

  for every value i in lower half of spectrum {
    matrix[i][REAL] = value[REAL];
    matrix[i][IMAGINARY] = value[IMAGINARY];
  }
  transposeMatrix;
}

cos_fft_max = getMaximumCosValue;
sin_fft_max = getMaximumSinValue;
fft_max = MAX(cos_fft_max, sin_fft_max);

factor = range / fft_max;
for every value i in matrix
matrix[i][REAL] = matrix[i][REAL] * factor;
matrix[i][IMAGINARY] = matrix[i][IMAGINARY] * factor;

```

Listing 11.4: CDM demultiplexer prototype: Pseudocode algorithm used to reconstruct images based on Fourier Synthesis.

```

for every row and column in image {
  recoveredValue = COSINES[row][REAL] + SINES[row][0];
  for every value in spectrum {
    // Base tone (over tones are ignored)
    recoveredValue += COSINES[row][value] * cos(2*PI*f);
    recoveredValue -= SINES[row][value] * sin(2*PI*f);
    recoveredImage[row][column] = recoveredValue;
  }
}

```

display so that it is possible to make sense of the perceived content. The screenshots in Figure 11.7 show four TV stations demultiplexed from a special video: (a) a newscast, (b) a sportscast, (c) a talk show, and (d) a movie.

Similar to the FDM prototype, the mobile application for the TDM method uses OpenGL ES 2.0 to rapidly process the live camera video feed. The algorithm shown in Listing 11.5 is used to determine whether the current camera image is a red synchronization frame. To speed up the analysis, the entire image is divided into 64 parts—8 rows and 8 columns—and only the four parts around the center of the image are considered in the analysis, i.e., 6.25% of the entire image. The algorithm computes the arithmetic means for each color, i.e., red, green, and blue, within these four parts. Eventually, the computed means are compared with the defined thresholds, which yields the final result whether the current frame is a red synchronization frame. To compensate for varying lighting conditions, e.g., outdoors vs. indoors, users may specify the applied thresholds via the user interface.

If the current frame is not a synchronization frame, i.e., it contains actual payload, the mobile application decides whether the current frame belongs to the selected information channel. The algorithm used for this is shown in Listing 11.6. To compensate for external environmental influences, e.g., varying timings of display panels, users may specify the accuracy threshold via the user interface. Higher threshold values may cause the mobile application to show frames that do not belong to the currently selected information channel, but to the one before or after. Lower threshold values avoid this, but may prevent the mobile application from recognizing the correct frames at all. Besides these thresholds, users may configure the number of available channels and set the frames per second via the user interface (not depicted here).

Listing 11.5: TDM demultiplexer prototype: Algorithm used to identify synchronization frames.

```
for (int i=0, int h=(height/8)*4; h < (height/8)*6; h++) {
  for (int w=(width/8)*4; w < (width/8)*6; w++) {
    i++;
    b = (b*(i-1) + redPixelValue (w, h)) / i;
    g = (g*(i-1) + greenPixelValue(w, h)) / i;
    r = (r*(i-1) + bluePixelValue (w, h)) / i;
  }
}

if (r > rThreshold && g < gThreshold && b < bThreshold) {
  return true;
}
return false;
```

Listing 11.6: TDM demultiplexer prototype: Algorithm used to extract the information channel selected by the user.

```
if (isSyncFrame) {
  expectedTime = CurrentTime() + (1.0/fps * informationChannel);
}
delta = fabs(CurrentTime() - expectedTime);

if (delta < accuracyThreshold) {
  showFrame();
} // else skip this frame.
```

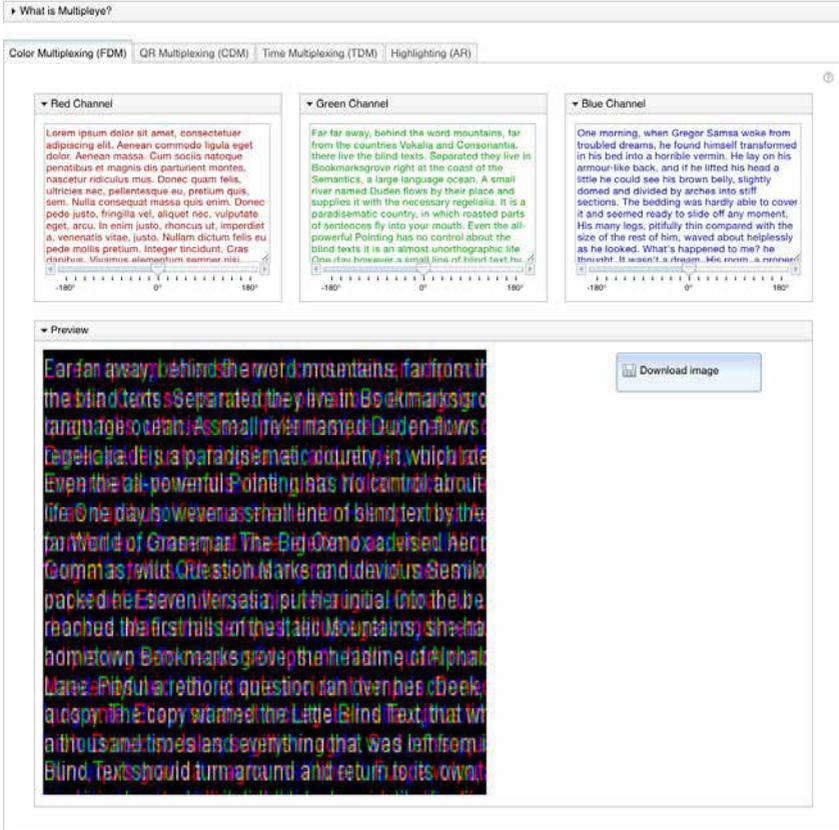


Figure 11.2.: FDM multiplexer prototype. This figure shows a screenshot of the web application. Users may define up to three individual texts as input images and download the multiplexed image. The image can then be imported into any third party software that is used to run a public display.

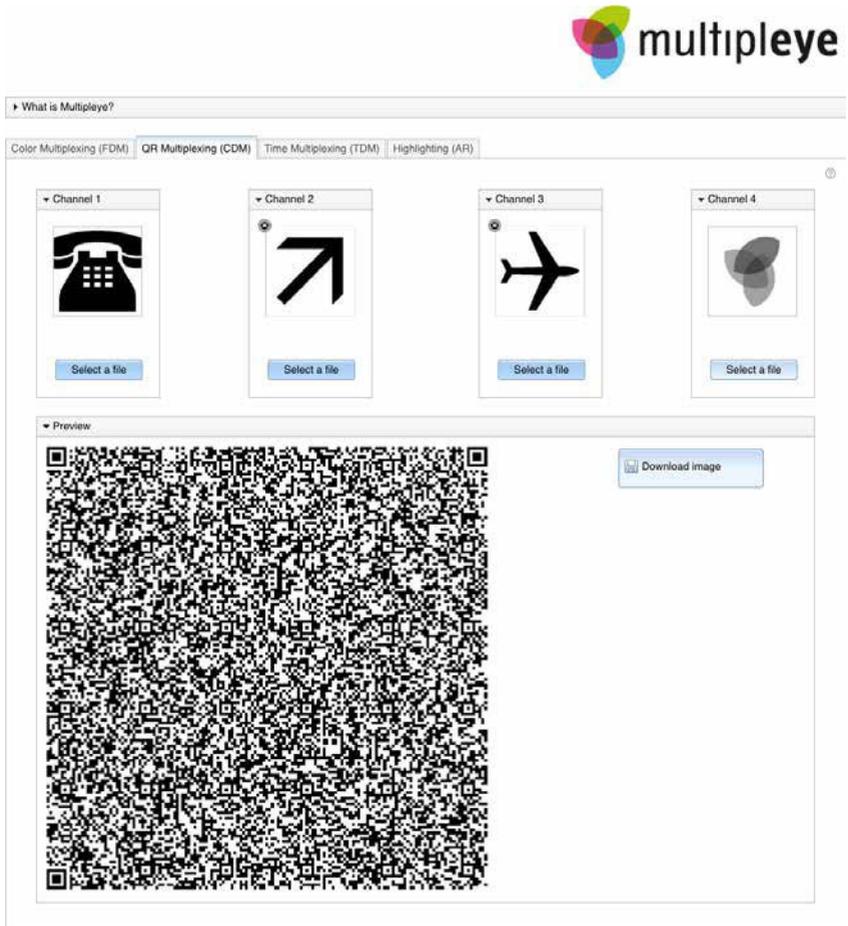


Figure 11.3.: CDM multiplexer prototype. This figure shows a screenshot of the web application. Users may define up to four individual input images and download the multiplexed image. The image can then be imported into any third party software that is used to run a public display.

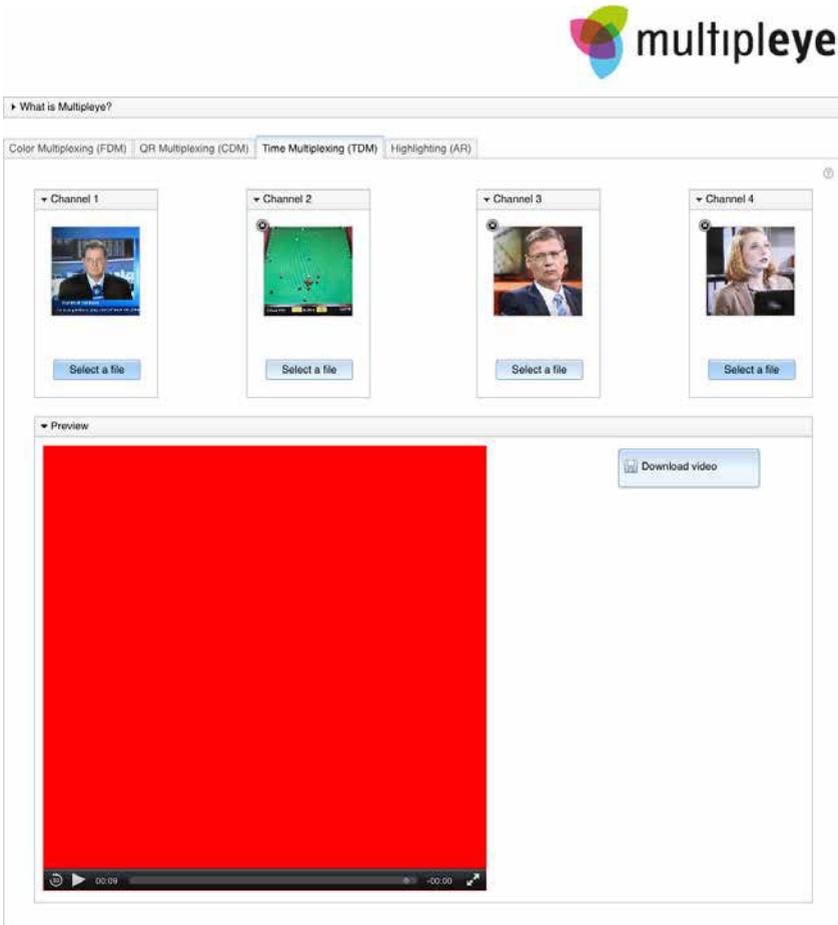


Figure 11.4.: TDM multiplexer prototype. This figure shows a screenshot of the web application. Users may define up to four individual input images and download the multiplexed video. The video can then be imported into any third party software that is used to run a public display. This screenshot shows the red video synchronization frame.

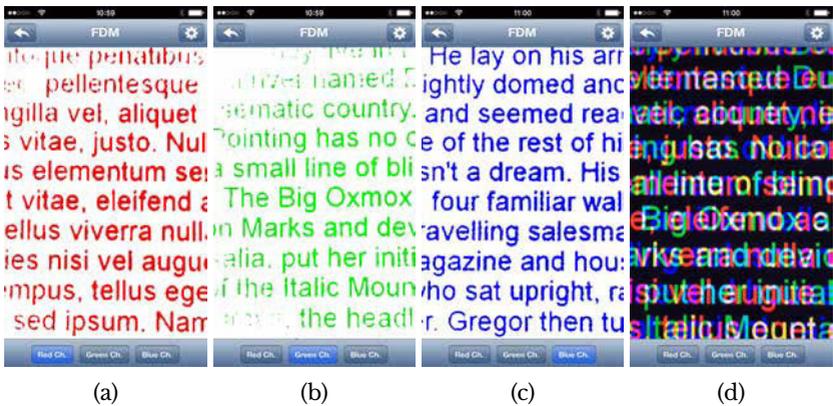


Figure 11.5.: FDM demultiplexer prototype. This figure shows screenshots of the mobile application. (a)–(c) Demultiplexed red, green, and blue information channel; (d) original camera image (no demultiplexing).

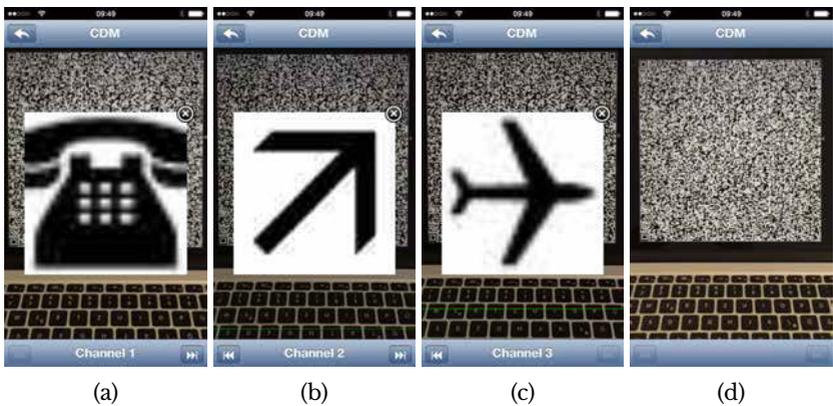


Figure 11.6.: CDM demultiplexer prototype. This figure shows screenshots of the mobile application. (a)–(c) Demultiplexed first, second, and third information channel; (d) original camera image (no demultiplexing).

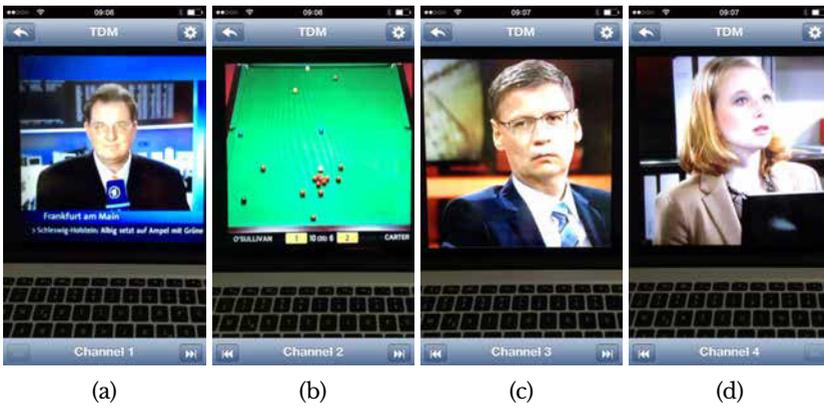


Figure 11.7.: TDM demultiplexer prototype. This figure shows screenshots of the mobile application. (a)–(d) Demultiplexed first, second, third, and fourth information channel.

11.2.2. Visual Highlighting

Similar to the prototypes for visual multiplexing presented above, the prototypical implementation for visual highlighting also consists of two components, i.e., a web application and a mobile application. Both components were integrated into the publicly available Multiplex prototype, see p. 265. The web application allows users to upload an arbitrary background image, on which visual highlights can be placed. The users may specify in which corner of the background image the QR/AR tag should appear. This option was considered necessary, as a particular background image could contain valuable information in either corner, which should not be covered by the QR/AR tag. Next, the users may define the coordinates, dimensions, and color of a rectangular highlight. A preview frame allows to align the highlight quickly and precisely on top of a specific visual information. Figure 11.8 shows a screenshot of the web application. In the screenshot, the user uploaded a background image that shows buildings and streets of a city as a high-angle shot. Furthermore, the users placed a red highlight on a building in the lower left corner of the background image. Based on this user input, the web application compiles an image that contains the background image and the specific QR/AR tag. The user may download the compiled image and use it in any third-party software to show it on a public display.

As the name implies, the QR/AR tag consists of two elements. The first element is a QR code that contains information about the highlights specified by the user. The QR code shown in Figure 11.8 holds the string “ARM0,-313.64,-137.2,98.64,79.2,ff0000.” The first three letters are used as *magic bytes*, that the mobile application uses to identify valid QR/AR tags. The next pair of numbers, i.e., -313.64 and -137.2, represents the X and Y coordinates of the highlight. The numbers are negative, as the origin of the coordinate system they relate to is located

in the middle of the QR/AR tag (axes facing the default directions, i.e., X from left to right and Y from bottom to top). The units does not directly relate to pixel coordinates of the chosen background image, because the used augmented reality framework (see below) applies a number of affine transformations, e.g., scaling. For the same reason, the numbers specified in the UI do not match the numbers in the QR code. The next pair of numbers, i.e., 98.64 and 79.2, relates to the width and height of the highlight. The same information about the coordinate system and the units applies here as well. The last string, i.e., ff0000, defines the color of the highlight, a deep red in this case.

The AR tag is a so called frame marker taken from the augmented reality framework used in the mobile application (see below). The web application uses the same web service to generate the QR codes as the CDM prototype described on pp. 267. Similar to the FDM prototype, the prototypical implementation of the visual highlighting approach uses the GD library to compile the QR/AR tag and embed it in the background image uploaded by the user, see pp. 266.

The mobile application also uses the ZBar library (see p. 267) to scan and decode the QR code contained in the previously described QR/AR tag. Once the coordinates of the highlight are extracted from the code, the mobile application uses the Vuforia augmented reality library¹³ to render the highlight within a virtual 3D space. Currently, the Multi-pleye prototype, i.e., the web application as well as the mobile application, supports one rectangular highlight, although—in theory—the number of arbitrarily shaped highlights is unlimited.

¹³<http://www.vuforia.com>, accessed: May 27, 2015.



Figure 11.8.: Visual highlighting prototype. This figure shows an edited screenshot of the web application. Users may upload a background image, define the position of the overlay (white border added to improve readability in monochrome printouts) as well as the QR/AR tag, and download the final image. The image can then be imported into any third party software that is used to run a public display.



Figure 11.9.: Visual highlighting prototype. This figure shows an edited screenshot of the mobile application. After scanning the QR/AR tag, the red visual highlight appears in the lower left corner (white border added to improve readability in monochrome printouts).

11.2.3. Visual Interaction

Subsection 10.2.3 introduced the design of the approach to visual interaction, which consists of two components. Both components were implemented as a prototype to evaluate the raw performance and general applicability of the approach. The prototypical implementation is named *Lichtblick*, a German word that refers to the main characteristic of the system to transfer data optically via light (Licht).

The first component, the public display, can run any software to display arbitrary content. *Lichtblick* only requires a small extra piece of software to run in the background (a *daemon*) that processes the camera stream. The daemon is implemented in C++ and uses OpenCV¹⁴. As soon as the daemon detects a particular light signal, it can trigger the corresponding action in two ways: (i) sending the received data to another process, e.g., the display software, via a TCP connection; or (ii) triggering a keystroke, mouse click, etc. The exposure time of the camera was set to a minimum so that the bit sequences of a light signal can be separated most precisely. Consequently, each camera image is quite dark for the human eye, see Figure 11.10a.

Once each camera image has been processed, i.e., gray-scale conversion, smoothing, binary black and white transformation, and noise reduction as depicted in Figure 11.10, the resulting image, shown in Figure 11.10e, can be analyzed by the daemon. White pixels represent possible light sources, i.e., the user's flashlight. The daemon then evaluates the presence of each light source in consecutive frames and deduces a stream of binary information, i.e., sequences of zeros and ones, e.g., "10101010 01000111 10101011." The string contains a start code (SC), a payload (p), and an end code (EC) (spaces included for readability only, see Table 11.1).

¹⁴<http://opencv.org>, accessed: May 28, 2015.

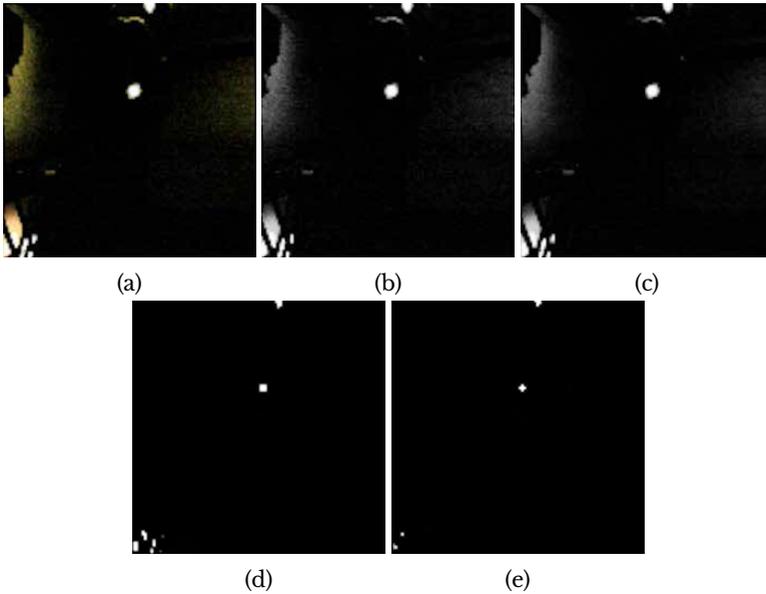


Figure 11.10.: Processing of the camera image performed in the visual interaction prototype. (a) Original camera image; (b) gray-scale conversion; (c) smoothing; (d) binary black and white transformation; (e) noise reduction, this is the final image used by the prototype.

Table 11.1.: Structure of a QR code with an example (concatenate table cells as strings from left to right).

	UI mode	SC	EC	Adaptive UI data	Frames per bit	Camera fps
Values	[A-Z][0-9]	8 bit	8 bit	text	integer	integer
Example	r	10101010	10101011	<pre><vbox> <button p="01000010" class="white"/> <button p="01010010" class="gray"/> <button p="01000111" class="black"/> </vbox></pre>	3	30

The second component, the smartphone application, is realized as an Android application. The application continuously scans for QR codes of a certain structure. Table 11.1 shows the structure (first row) as used in the prototype and also provides a concrete example (second row). However, the actual structure is implementation specific and not relevant to the general concept of optical interaction; it is thus not discussed in detail here.

The UI mode defines how the public display should interpret the light signals emitted by the smartphone. The display could, e.g., move the mouse cursor according to the movement of the smartphone if the UI mode is set to “c” (cursor); the UI mode “r” (remote) tells the prototype to handle light signals like keystrokes as emitted by a TV remote control, for example. The start code (SC) and the end code (EC) should be defined so that their corresponding binary representations may not occur in the payloads (p) of the light signals. The adaptive UI data contains information about how the user interface on the smartphone should look like (here defined in XUL). The frames per bit specifies for how many frames a light source has to light up to be considered as turned on. The camera fps defines the frame rate of the camera. Both, frames per bit and camera fps have been added to provide more flexibility with varying hardware setups and installation environments.

Once the application scanned a QR code, it deduces the required communication parameters, constructs the adaptive user interface (see Figure 10.12e, for example) and continues to scan for updated codes. A particular challenge while implementing the prototype was caused by the fragmentation of Android devices and software versions. Various manufacturers install different versions of the Android operating system on their smartphones. Consequently, there is a broad range of different hardware components with varying characteristics. Controlling the flashlight of the camera precisely turned out to be an is-

sue: Some devices allow applications to control the status of the flashlight directly, i.e., with minimal delay; other devices, however, do not provide this direct access, so that apps have to use the camera API provided by Android. Although the latter may be a more general approach, it may cause problems with respect to performance, as the additional software abstraction level appears to introduce delays. Subsection 12.2.3 reports on observations about this phenomenon.

11.3. Process Integration

As with the other approaches presented in this thesis, the scientific contribution of a process integration (C3) was realized as a prototypical implementation for two reasons: Firstly, to provide tangible contributions and secondly, to allow for a well-founded evaluation. The remainder of this section thus presents the prototypes of the IPED Toolkit and the Immersive Video Environment.

11.3.1. Immersive Public Display Evaluation and Design Toolkit

The IPED Toolkit is an initial prototypical implementation of the approach proposed in Section 10.3. This subsection describes the toolkit, its core elements, its design and evaluation processes, and the current implementation in more depth. Figure 11.11 provides an overview of the architecture and the components of the toolkit.

The *core* component maintains the state transition graph that defines the simulated world and manages all information that is attached to the graph. In particular, it contains the (video) footage that represents

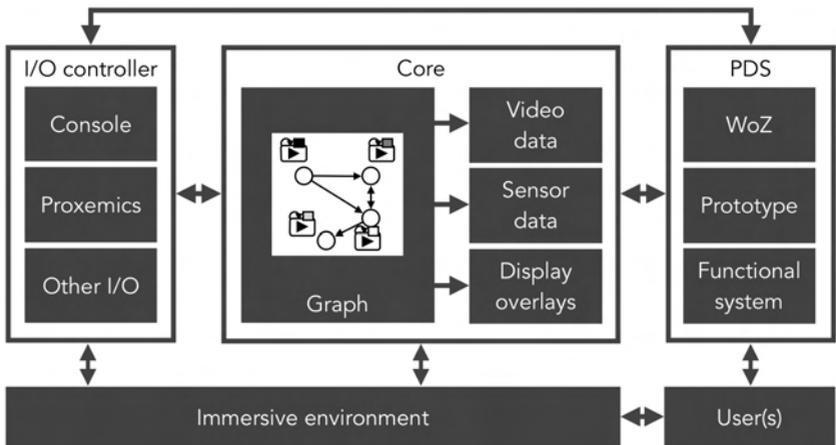


Figure 11.11.: Architecture of the IPED Toolkit. The I/O controller wraps different means of interaction; the core contains the graph and data defining the simulated world; the different public display systems (PDS) control the behavior of the system simulated in the immersive environment.

the real-world locations and their states as well as the positions of virtual public displays. Optionally, it can store additional sensor data that has been recorded at particular locations. Finally, it maintains the user's position in the simulated world.

The *PDS* (public display system) component controls the contents of the virtual public displays that are embedded in the simulated world and it reacts to any interaction. Content can be generated in different ways, e.g., via a Wizard-of-Oz (WoZ) approach, i.e., that a human “wizard” selects appropriate content in response to what a user does in order to simulate a functional system. Alternatively, prototypes of different degrees of sophistication can (semi-) automatically generate the screen contents. Equally, it is possible to use a fully functional version of a public display system to generate screen contents and react to the user's input. Content can be represented as, e.g., images, videos, web pages, or screen casts (see p. 295). This way, it can be displayed via standard web browsers within the overlays that represent virtual public display systems.

The *I/O controller* encapsulates different means of interaction between the simulated world and the public display systems. It incorporates a console application that can be used to configure the simulated worlds and to control the user's movements within those worlds. Additionally, the I/O controller includes a simple mapping API that different interaction mechanisms may use to control various aspects of the simulated worlds, such as moving between locations and states or triggering specific actions, for example. The latter capability allows for interaction with public displays via proxemics: The user's location in the lab can be translated to the corresponding location in the simulated world and thus trigger a reaction of the public display in turn.

Users experience the system via an Immersive (Video) Environment, see Figure 11.12 and Subsection 11.3.2. They can interact with the system in a variety of ways. What they see is determined by their position within the graph, which is managed by the core, and by the content of simulated displays, generated by the PDS component. The overall architecture thus decouples the simulated world, the display placement and content, as well as the way in which people interact with a public display system. This structure provides a high degree of flexibility with respect to designing such a system and modifying aspects of it. It also facilitates the re-usability of components and the reuse of simulated environments.



Figure 11.12.: Panoramic video footage presented in the Immersive Video Environment. Three large screens show high-definition video footage of the old train station in the city of Münster (recorded in summer 2014). The virtual public display in the center (left to the traffic light) can be positioned in the 3D space as indicated by the green lines of the correlated coordinate system.

The prototype of the IPED Toolkit is realized as a web application. This platform was chosen for a number of reasons: Firstly, web development and corresponding technologies made a lot of progress within recent years. Browsers became more and more powerful and a large (open source) community contributes a plethora of libraries and frameworks that the prototype may profit from. Secondly, the deployment of web-based applications is straightforward in comparison to other approaches, e.g., Java, Grails, or Ruby. As the prototype is publicly available on GitHub¹⁵, researchers or designers can simply download the sources, launch the built-in web and database servers, and start using the toolkit. The underlying technology, i.e., Node.js¹⁶, keeps track of all dependencies and downloads missing packages automatically. Finally, as web applications became very popular and common within recent years, many developers are likely familiar with the employed technologies and design patterns. They may thus be able to customize the prototype to suit their individual needs.

To develop the prototype in a contemporary manner and to provide a modern user experience, the project is built on a number of well-established projects: The web server is implemented in Node.js, while the transition graph is maintained in a Neo4j¹⁷ database. In contrast to other databases, such as MySQL or PostgreSQL, Neo4j stores data in a graph-based model itself. Thus, it provides sensible methods to manipulate and traverse graph-based structures, i.e., the transition graph of the the IPED Toolkit.

Furthermore, a professional technology stack was used to ensure a constant level of code quality. For example, the Grunt¹⁸ task handler first checks all source code files for syntactical errors, then “beautifies”

¹⁵<https://github.com/sitcomlab/IPED-Toolkit>, accessed: June 2, 2015.

¹⁶<https://nodejs.org>, accessed: June 2, 2015.

¹⁷<http://neo4j.com>, accessed: June 2, 2015.

¹⁸<http://gruntjs.com>, accessed: June 2, 2015.

the sources, and finally merges all individual files into one file to reduce data transmission times. A continuous integration system, i.e., Jenkins¹⁹, automatically updates both production and development servers. Overall, established design patterns were applied, in order to guarantee reliable and readable source code that other researchers and designers may easily adapt. To enforce these design patterns, widespread frameworks such as Backbone.js²⁰ and Require.js²¹ were used. Finally, the prototypical implementation of the IPED Toolkit is well-documented, for example, via automatically generated API specifications. Tools such as Apiary²² or JSDoc²³ are used to keep the documentation up to date.

The prototype of the IPED Toolkit consists of two web pages, each addressing a specific task. The general concept is aligned with common content management systems, that provide dedicated views for individual users: Editors use the backend to create and manage actual contents, i.e., locations, relations, and overlays; users may experience the contents created by the editors in the *frontend*. Each component is presented in more detail below.

¹⁹<https://jenkins-ci.org>, accessed: June 2, 2015.

²⁰<http://backbonejs.org>, accessed: June 2, 2015.

²¹<http://requirejs.org>, accessed: June 2, 2015.

²²<https://apiary.io>, accessed: June 2, 2015.

²³<https://github.com/jsdoc3/jsdoc>, accessed: June 2, 2015.

Backend

Editors, e.g., designers or researchers, can use the backend to manipulate the transition graph. The main component of the backend is a Leaflet²⁴ map, that visualizes all available situations and their correlations, i.e., links between them. The backend provides a contemporary user interface based on technologies such as Bootstrap²⁵ and jQuery²⁶. This user interface allows to create, edit, and delete situations, video footages, and overlays. Figure 11.13 is a screenshot of the backend that shows information about a particular situation called “Windthorststraße.” This situation is linked to three other situations (that means that users can virtually walk to these locations), is based on one video footage, and features one overlay (a public display in this case, see Figure 11.15).

As explained in Subection 10.3.1 on pp. 249, the edges of the transition graph represent links between situations. In most cases, a link corresponds to an actual path in the real world, that users may take to get from one location to another. However, links can also be used to model different states of a particular situation, e.g., daytime vs. nighttime. To create links between two situations, editors may right click on the particular two locations and select the direction of the new link. The IPED Toolkit supports both, uni-directional as well as bi-directional links. Figure 11.14 is a screenshot that shows how an editor creates a link from the situation with the arrow facing upwards to the topmost situation.

Another feature of the backend is the ability to create overlays, e.g., virtual public displays, on top of the recorded video footage, see Figure 11.15. Editors can use the coordinate system indicated by the green

²⁴<http://leafletjs.com>, accessed: June 2, 2015.

²⁵<http://getbootstrap.com>, accessed: June 2, 2015.

²⁶<https://jquery.com>, accessed: June 2, 2015.

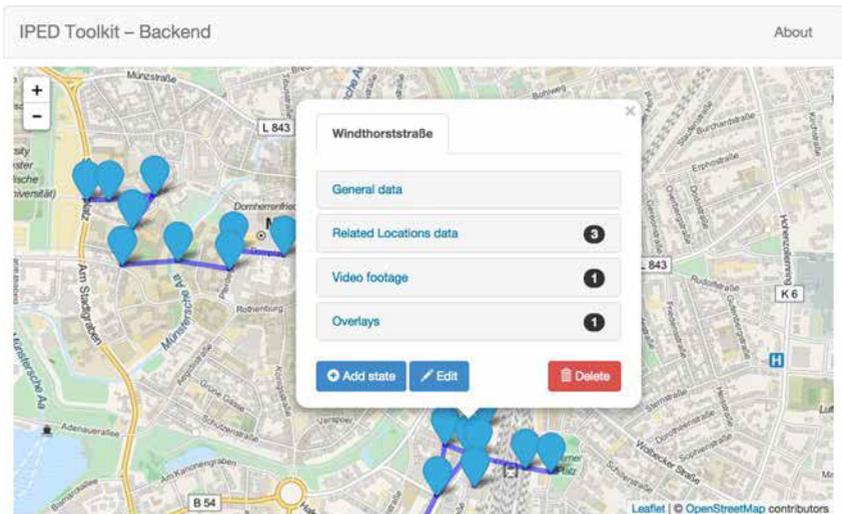


Figure 11.13.: Screenshot showing the backend of the IPED Toolkit with the underlying map, a number of situations (blue markers), links between these situations (blue lines), and an open details pane on top of that map.

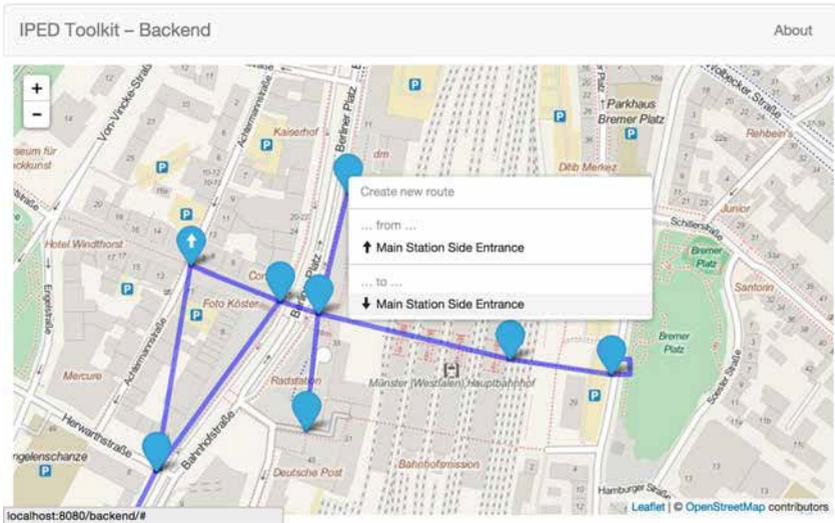


Figure 11.14.: Screenshot showing how the backend of the IPED Toolkit can be used to create links between situations.

lines laid on top of the background video to position the overlay within a 3D space. To have the overlay blend in as seamlessly as possible, editors can translate, rotate, and scale the overlay in each of the three dimensions. The 3D manipulation is based on `three.js`²⁷, a powerful JavaScript library.

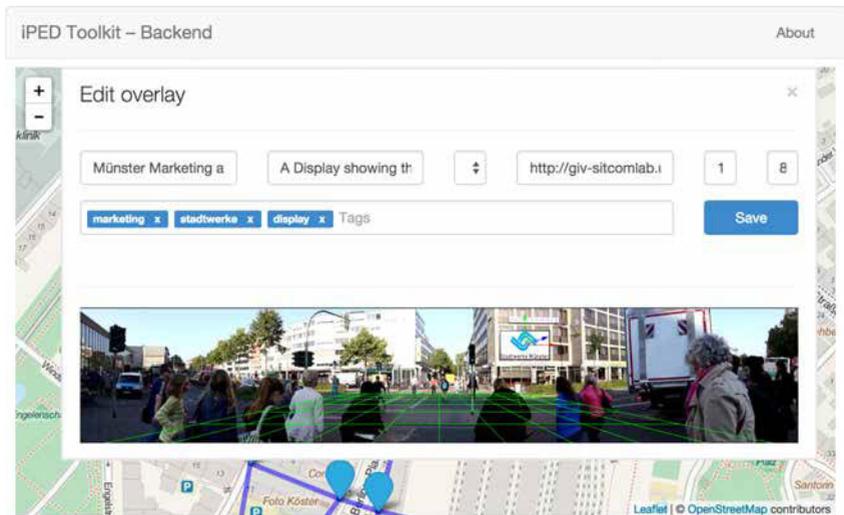


Figure 11.15.: Screenshot of the IPED Toolkit showing how the backend can be used to create overlays on top of the recorded video footage; here, the overlay is a public display showing an advertisement for the public utility company (Stadtwerke Münster).

²⁷<http://threejs.org>, accessed: June 2, 2015.

Frontend

While the backend is intended to be used by content editors, the frontend presents the simulated environment to the user. One particular advantage of the web-based approach is that the frontend can be displayed on different devices in various resolutions. Figure 11.12, for example, shows the frontend running as a fullscreen application. Most browsers can easily be put into fullscreen mode by pressing a dedicated key, e.g., F11 in Firefox. Since the frontend does not rely on a particular display method or resolution, it can be used in multiple application scenarios. For example, users may experience a public display system in a large-scale Immersive Video Environment, as shown in Figure 11.12, or perceive a smaller experience on a laptop display. This flexibility allows researchers and designers to present their simulations outside their labs at, e.g., conferences or exhibitions.

Another feature of the frontend is that it allows third-party content to be projected into the simulated environment, as depicted in Figure 11.16. The underlying technology is webRTC²⁸, which allows for real-time communications (RTC) between browsers. In the application scenario depicted in Figure 11.16, the laptop runs a prototypical implementation of a software that is designed to drive a public display. The screen of that laptop is captured and transferred via webRTC, so that it can be displayed at a specific location within the frontend. As soon as the content shown on the laptop changes, these changes are immediately reflected in the simulation. Since the screen capturing process is technology-agnostic, any software can be used to drive the simulated content: still imagery, videos, websites, or native applications based on Java or C are just some examples.

²⁸<http://www.webrtc.org>, accessed: June 2, 2015.

Another feature builds on top of this: interactive sketching. Instead of showing the prototype with three buttons, the laptop shown in Figure 11.16b could run an arbitrary image processing or sketching tool, such as Inkscape²⁹ or Gimp³⁰. The user can now draw arbitrary sketches on a green canvas. The green parts of the background will be removed in the frontend, so that only the drawn sketches remain visible. This technique is commonly known as green-screen compositing or chroma-keying. The used JavaScript library Seriously.js³¹ provides a fast and robust implementation for this.

Remote Control

The remote control is an additional web page to control the simulated environment. The user interface is designed to be accessible on mobile devices, e.g., tablets or smartphones, as well as on desktop computers, see Figure 11.17. When conducting user studies, for instance, the remote control allows experimenters to select locations or routes represented by nodes and edges in the transition graph. As soon as the experimenter selects the desired item, the simulation shown in the frontend changes accordingly. The communication between the remote control and the frontend is realized via WebSockets³².

More sophisticated approaches to control the simulated environment are, for example, voice control or gestures. The first approach was developed and evaluated in the context of a bachelor thesis (T6, see p. 443). Besides showing the general feasibility, the thesis contributed a list of the most popular expressions people would use to control the

²⁹<https://inkscape.org/de>, accessed: June 2, 2015.

³⁰<http://www.gimp.org>, accessed: June 2, 2015.

³¹<http://seriouslyjs.org>, accessed: June 2, 2015.

³²<http://tools.ietf.org/html/rfc6455>, accessed: June 2, 2015.



(a)



(b)

Figure 11.16.: Projection of third party screen content into the simulated environment. (a) Screenshot of the simulated environment with a public display overlay showing the content of a remote laptop; (b) the remote laptop running the actual prototype of a public display software.

simulation via voice commands. Furthermore, the evaluation indicated that using the voice control does not impose a significant mental or physical workload on the users. In contrast, being able to control the simulation with words only—instead of pressing physical or virtual buttons—helps to increase user’s the feeling of being immersed.

Another thesis analyzed the latter approach (T4, see p. 443). Specifically, the thesis compared dynamic gestures with poses used to control Immersive Video Environments. Besides serving as a proof of concept, the results indicate that gestures let users feel more immersed than poses. However, users also claimed that the usability of the system appeared to be “better” when using poses. Ostensibly, this result was related to the capability of the system to recognize and differentiate gestures quickly and reliably. Future studies may provide further insights into this.

Example Scenario

The IPED Toolkit can be applied to different stages of the development process of a public display system as illustrated in the following example scenario, cf. Figure 10.13: A network of public displays is to be installed at a large hospital in order to replace static signage. One key function could be to provide individuals, e.g., patients, visitors, or employees, with personalized directions.

Using the IPED Toolkit, the first step would be to identify the sites in which the system should be installed, e.g., where current static signage is mounted. Next, the required video footage needs to be acquired. One option would be to reuse existing footage that was previously recorded, for example, via a public repository. However, since the hospital is a very specific scenario, it is unlikely that suitable footage

IPED Toolkit Remote

The screenshot displays the IPED Toolkit Remote web application interface. It features three main sections:

- Select a start location:** A blue header bar.
- Select a route:** A blue header bar above a white container with three buttons: "Bahnhofsstraße", "Windthorststraße / Achtermannstraße", and "Main Station".
- Settings:** A light gray header bar above a white container with two settings:
 - Show Overlays:** A toggle switch currently set to "ON".
 - Voice Control System:** A toggle switch currently set to "OFF".

Figure 11.17.: Screenshot of the remote control web application. Users may select a start location or a route and control some settings of the simulated environment, for example, disabling overlays.

exists. Thus, the required footage would have to be recorded in situ. Usually, a few minutes per site are sufficient to create seamless loops.

After recording videos at the identified locations, the footage has to be processed, e.g., for seamless looping. Most often, it is only necessary to crop videos so that the beginning and the end of the clip are visually similar. Then the graph that links the different sites and recordings has to be constructed. The nodes of the graph correspond to the decision points identified earlier. They are linked to the corresponding video footage. The edges between two nodes represent physical or logical connections between decision points, i.e., adjacent locations are most likely connected by an edge. Eventually, the graph contains the required information for simulating the deployment site. It can now be easily reused in other contexts as well, e.g., to develop a public display system managing waiting times in the same hospital or to develop an application in a more generic hospital scenario.

Once these steps are completed, the IPED Toolkit can be used to test and discuss the placement, shape, or size of public displays at different locations by overlaying the corresponding designs over the recorded footage. Using the toolkit greatly simplifies designing and prototyping at this stage, as designers do not need to be at the actual deployment sites. They can rather place (virtual) displays freely and rapidly without physical efforts. The placement and configuration of each public display is stored with the corresponding node. This data can be reused in different contexts as well, since the content of the displays is controlled and generated independently.

As the development of the system progresses, designers and users can use the toolkit to test, inspect, and discuss the system at various stages. For example, at early stages, a static mock-up of the proposed dynamic signage system could be used to assess whether the interface design fits the targeted installation site. Once a functional dynamic

prototype exists, it can be connected to the display overlays within the simulation to replace the static mock-ups for further analysis.

If there are multiple recordings for one location, e.g., representing busy vs. quiet office hours, these recordings can also contribute to the evaluation of the system prior to its actual deployment. In the hospital scenario, it would be possible to test whether the content of the public displays is unambiguous and thus suitable to guide visitors through the building. This could be done by assessing the user's performance in terms of task completion times, i.e., how long it takes them to determine which way they need to go at each display, or error rates, i.e., how often they take the wrong direction at a display.

Finally, the IPED Toolkit also supports the integration of different means of interaction. In the hospital scenario, designers might want to test an interface based on proxemics [133]: If a person steps closer to a display, its content could change to provide more detailed directions for that person. Sensing the relative position of a person to a screen in the lab can be achieved easily, e.g., by using depth cameras. These sensors could then be connected to the prototype of the public display system, which in turn adapts the screen content that is shown in the simulation accordingly.

Once connected, a sensor is available at all simulated locations without the need to physically deploy it multiple times—which would be necessary for field trials. Compared to designing, testing, and evaluating public display systems in the real world, the toolkit can thus reduce effort at several stages. It simplifies the integration of sensors and allows for covering large areas without the need for extensive hardware deployments. At the same time, the toolkit can simulate and control context realistically. It also facilitates the reuse of simulated installation sites and simulated displays. While the approach can thus complement existing methods well, it is also subject to some limitations.

For example, locomotion is only possible within very narrow limits, as is the simulation of physical aspects, e.g., temperature, texture of surfaces, or precipitation. Furthermore, appropriation and user acceptance can only be assessed in a limited way. The interaction with passersby and other users is subject to constraints as well.

Summary

In its current version, the IPED Toolkit can simulate the following aspects that are relevant in the development of public display systems: In terms of the actual display, the toolkit can incorporate the location and orientation at a specific site (first challenge, see Section 9.1), the size, shape, and form factor (second challenge, see Section 9.2), as well as the content (fifth–eighth challenge, see Sections 9.5–9.8). Contextual factors that can be simulated to some degree include visual environmental properties (third challenge, see Section 9.3), ambient noise, time of the day, and—in a limited way—passersby and bystanders, as recorded in the footage. Some events, such as an arriving bus or an opening door, can also be simulated (fourth challenge, see Section 9.4). However, quite a few contextual factors cannot be simulated easily, for example, temperature, precipitation, or the negotiation of a crowded place. The IPED Toolkit facilitates the interaction with public displays via mobile devices (fifth challenge, see Section 9.5), gestures, and also via proxemic interaction—though the range of locomotion is limited. In addition, it supports the rapid development of multi-display networks (sixth challenge, see Section 9.6) within the limitations outlined above.

11.3.2. Immersive Video Environment

The approach towards a process integration as presented in this thesis uses an Immersive Video Environment to simulate public display systems and their surroundings, see Figure 11.12. The IVE presented in this subsection is inspired by the system presented by Singh et al. [209]. Panoramic video footage recorded at real-world locations is played back in order to immerse users into these settings. Three back-projection screens (219 cm wide, 164 cm high) are arranged in a semi-circular manner, spanning a viewing angle of about 114 °. Each screen is driven by a short-range back-projector. The resolution of each projector is 1920 x 1080 pixels, resulting in an overall resolution of 5760 x 1080 pixels. The corresponding footage is recorded using three standard DSLR cameras (Canon EOS 550D) mounted on a custom-made tripod. The angles between the cameras are adjusted to match the angles of the three IVE screens. To minimize visible seams between the three videos, the footage is post-processed using standard video editing software. A high quality surround audio recorder (Zoom H2n) is used to capture ambient sounds in decent quality. Both, audio and video material, is then played back on a single desktop PC, equipped with a graphics card that can drive multiple displays.

Mirror Image Avatar

Subsection 10.3.2 introduced the approach of using a mirror image of the user as an avatar embedded in the Immersive Video Environment. The main goals of the approach were to enable intuitive interaction with video environments and to immerse people as much as possible into the real world scene being shown. The aim was thus to enable users to perform various actions in the simulated environment while providing them with a strong feeling of presence. The basic idea is

to create a realtime mirror image of the user and overlay it over the video footage. Using a simple layer-based depth model and a small set of gestures, users can place their mirror avatar inside the depicted real-world scene and interact with the environment via the avatar, e.g., to place virtual objects. The prototypical implementation to this approach is presented below.

A live mirror image of the user constitutes the focal point of interaction, see Figure 11.18a–d. The background behind the user, which is also captured by the camera pointed at the user, is eliminated in real-time, e.g., using standard chroma-keying, see p. 296. The cut-out mirror image serves as an avatar or proxy for the user: Its location in 3D space defines where interaction can take place.

Specifically, the avatar defines the currently selected depth layer. Since the user and the PC (or the depth camera) know how tall the user is, the actual size of the avatar as depicted on screen defines its depth position inside the video footage. The size of the avatar is used as a depth cue to inform the user about the layers and objects that can be selected. For example, by placing the avatar on the third layer, users can interact with objects located on that layer and can inject additional virtual objects in that layer. The remainder of this subsection describes the gestures that can be used to control the avatar and the simulation.

Avatar Gestures. In order to move the avatar within the 3D space defined by the video footage and the layer model, users can perform a number of gestures. When one foot is placed in front of the other, users can move their avatar along the X-, Y-, and Z-axes. To control the movement in X- and Y-directions, users use a one-handed gesture: By extending one arm in one of the cardinal directions, users can specify in which direction the avatar should move, see Figure 11.19a.

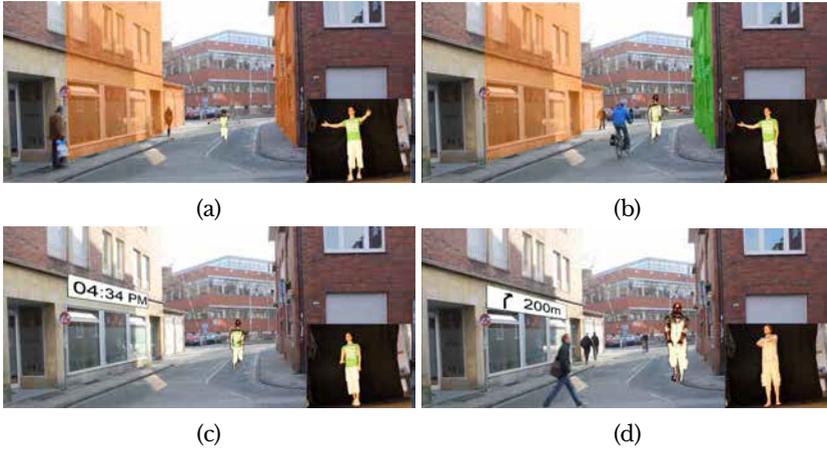


Figure 11.18.: Mirror image avatar example application. (a) Moving the avatar using gestures; objects on the current depth layer are highlighted; (b) pointing in the direction of objects on the current depth layer selects them; (c) using gestures to move and scale virtual objects, here a public display showing the current time; (d) moving the avatar to experience the virtual environment, i.e., the virtual public display reacts to the avatar.

For example, extending the arm to the left moves the avatar in that direction. Movement continues while the arm is extended; movement stops when users bring their arm close to their bodies or when they perform a different gesture.

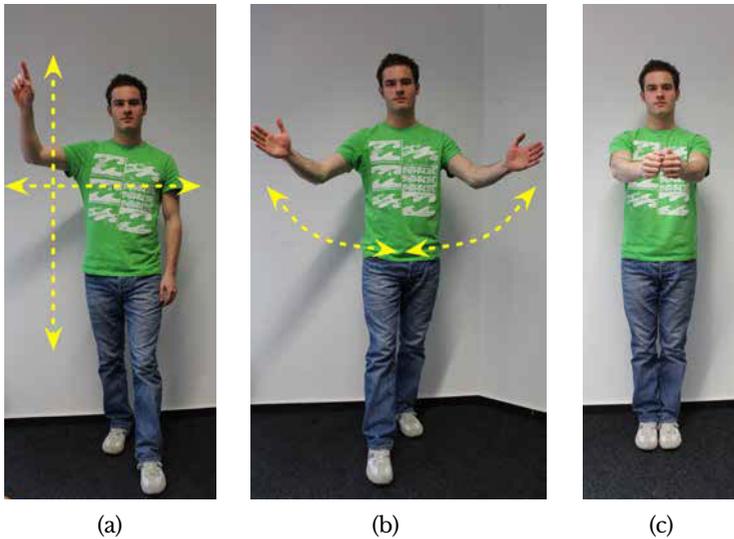


Figure 11.19.: Key gestures used to control the avatar, interact with the video scene, and to manipulate virtual objects. (a) Moving gestures; (b) scaling gesture; (c) switching gesture.

A two-handed gesture controls movement along the Z-axis, i.e., the depth. Putting both hands closely together in front of the body shrinks the avatar. This corresponds to moving it deeper into the image, i.e., it increases its distance to the camera. By spreading both arms, users can increase the size of the avatar and thus decrease its distance to the camera, see Figure 11.19b. The scaling stops when users either return their arms to a relaxed position, or when they perform another gesture. The size of the avatar determines which depth layer is selected: The depth information specified for each layer and the actual height of the user enables the system to compute the best match, i.e., to find the layer which corresponds best to the current size of the avatar.

An initial user study evaluated these gestures by comparing them to an alternative set (T4, see p. 443). That second set allowed to control movement by walking in place while orienting one's body in the target direction. The key findings are that participants were largely successful in navigating the avatar to the target locations using either set. A frequent observation was that people stop when their avatar reached a street in order to avoid collisions with cars—this indicates a high degree of immersion. The gesture set depicted in Figure 11.19 was rated more favorably than the comparison set. Overall, 70% of the participants preferred the static gesture shown in Figure 11.19.

Object Manipulation Gestures. In addition to moving the mirror image avatar, a set of gestures was defined to select objects depicted in the video footage, e.g., buildings or signs, to inject virtual objects, such as public displays, new buildings, or audio sources, as well to move and scale those objects. In addition, gestures to select content to inject into the scene or other content-related activities can be defined. For example, the prototypical implementation included gestures to select items from a list of options in order to add them to the scene.

In order to select an object in the video footage, users first need to place their avatar on the corresponding layer. Users then select an object by simply pointing in the direction of the object. The gestures to place an object in 3D space are the same ones as those used for moving the mirror image avatar. Users can switch between moving their avatar and moving objects by putting their hand together, extending their arms in front of them and then maintaining this pose for a short time, see Figure 11.19c. Visual feedback indicates the switch from one set of actions to another.

Experience Gestures. Once the video footage has been augmented with a number of virtual objects, people can experience the new scenario in the following way: They can interact with the objects augmenting the video scene by moving their mirror image avatar through the 3D space defined by the video footage and the layer model. The layer model provides means to, for example, measure the distance to a public display and realize proxemic interaction [24]. This set of gestures is similar to the avatar gestures presented above.

12

Evaluation

The previous chapter presented the prototypical implementations of the approaches presented in Chapter 10, i.e., the privacy threat model, the three novel countermeasures, and the proposed process integration. The twofold rationale for these prototypes was to provide tangible scientific contributions, but also—more importantly—to use them as the basis for a well-founded evaluation. This chapter reports on the evaluation of the proposed approaches and their corresponding prototypes. Section 12.1 presents the results of a user study with regard to the privacy threat model (C1). Section 12.2 is divided into three subsections, each concerned with one of the three proposed countermeasures (C2), i.e., visual multiplexing, visual highlighting, and visual interaction. Finally, Section 12.3 focuses on the evaluation of the suggested process integration (C3).

All sections assess the suitability of each approach and report on qualitative results obtained from user studies or interviews. Besides this evaluation, however, the sections also aim to assess the technical feasibility of the proposed approaches, since this thesis strives to provide concrete scientific contributions.

12.1. Privacy Threat Model

The first scientific contribution (C1) of this thesis is a privacy threat model for interactive public displays, see Section 10.1. The prototypical implementation of that model was introduced in Section 11.1. The purpose of the prototype is twofold: Firstly, it enables other researchers and designers to directly apply the proposed privacy threat model to their projects; secondly, it turns the theoretical model into a concrete and tangible tool, that may be used to demonstrate and evaluate the underlying concept. This section presents the results of a qualitative as well as a quantitative evaluation of C1. The qualitative evaluation is based on a semi-structured interview, while the quantitative evaluation used a questionnaire.

The semi-structured interview was disseminated via e-mail to four renowned experts (E1–E4) in the field of public displays (E1: 2 journal articles and 9 conference papers; E2: 2 journal articles, 11 conference papers, and a dissertation; E4: 33 conference papers, 5 journal articles, and 2 book chapters) and privacy research (E3: 34 conference papers, and a dissertation). The experts were briefly introduced to the topic and then asked to look at the prototype of the privacy threat model. They were then requested to answer three questions about (i) the extent of the model, (ii) its usefulness to designers, and (iii) general comments or remarks regarding the model. The entire letter including the three questions is included in the appendix on pp. 451.

In addition to the expert interviews, a paper-based questionnaire was handed to twelve students (S1–S12), who used the privacy threat model in a seminar on location-based privacy. The seminar was taught jointly with the University of Minneapolis. Students and teachers from the cities of Minneapolis and Münster collaborated throughout an entire semester on multiple location-based services. For example, one sys-

tem navigates users between two locations while avoid being tracked by surveillance cameras (CCTV). The privacy threat model and its prototype should be used to design these location-based systems.

The questionnaire contained five questions: Similar to the questions presented to the experts, the first question was about the extent of the model and the second question asked for the usefulness to designers. The third and fifth question tried to evaluate whether, and why, students would have the feeling of an increased privacy if any of the three countermeasures (C2, see Section 10.2) was applied. The fourth question comprised a *NASA TLX* questionnaire to assess the perceived workload and usability of the threat model. Table 12.1 summarizes the corresponding results. Finally, there was space for general comments.

In addition to the questionnaire, the students were asked to use the prototype of the privacy threat model to create models for the systems that they were to develop in the course of the seminar. As the students worked in pairs, there are six resulting privacy threat models, which are shown in Figure 12.1. The models were edited afterwards to improve readability in monochrome printouts. They are described in the remainder of this section and Section 14.1 discusses the outcomes. Overall, the answers provided by the experts were positive with regard to the estimated extent of the threat model. Yet, the comments also suggest ways to improve the model and indicate further research directions, for example:

In my opinion, the model is quite comprehensive and covers major potential attacks. The examples, e.g. for the threats, are very helpful. Some of the threat categories are slightly overlapping (e.g. in tampering/information disclosure – I’m not quite sure whether there is partly a mixture of action and result). One aspect that seems to be missing is the combination or cooperation of threats or agents. Can I model an attack where two malicious persons are involved or an attack that combines two or more weaknesses? I guess that depends on the interpretation of child nodes. What does the link between an agent and a threat exactly mean (origin, exposed to?)? In case of a benign or a malicious user, the meaning seems to be different. —*Expert 1*

The extent of the model seems reasonable. The most common threads [sic] are included. However, it would be great if additional own threads [sic] could be added since public display settings often have unique characteristics that may introduce additional threads [sic]. A further aspect that I consider as very important is the spatial arrangement of the setting and the interaction. I would suggest to add such a spatial component to the model. The orientation, distance to other users and the display, as well as viewing angles are important when displaying content on a public display. —*Expert 2*

[...] It does look comprehensive. —*Expert 4*

I think it’s okay.

—*Expert 3 (English translation by the author of this thesis)*

The remarks on how the threat model and its prototype could support the design of privacy-preserving public displays were also positive. Again, the experts provided valuable suggestions about how to improve the privacy threat model in future work, too:

The model definitely supports the “creators” of public display installations by making the potential attacks visible and raising awareness. One might argue, that such a threat model looks the same for each interactive public display, i.e. there is only one threat model (with potential attacks, involved actors, etc.) for all displays. Which factors affect the model for a specific installation? The type of application? The available input possibilities [sic]? —*Expert 1*

I think the threat model can help designers of public displays to design privacy-preserving systems in terms of providing means to model threats in advance. It would be a further benefit for designers, if they could model their system with the tool, and the threat model analyzes the the model and notifies the designer of possible threats that need to be addressed in the design (and this can then also be done within the tool). —*Expert 2*

I also liked that the tooltips [...] provided easily understandable examples. They definitely help to raise the user’s understanding.

—*Expert 3 (English translation by the author of this thesis)*

The fourth expert, however, articulated some concerns with regard to the usability of the prototype and provided constructive advice:

[...] I always get this feeling that I am not sure if the results I get with the model would actually apply to whatever system I may be designing.

The second concern is that it feels awkward that I have to create the entire model. It seems I am building the wall and the cannons to destroy it. How can I ever know that I am being comprehensive?

[...] Putting my self [sic] in the shoes of someone creating a display system that wants to use a tool that would help to uncover system vulnerabilities, I would expect more. I would probably expect to describe the key points of the system, perhaps in a wizard-like questionnaire, and then be pointed out with potential vulnerabilities. I would then describe existing or planned countermeasures until the model was “secure.” Basically, my point is that if I do not have expertise on threat models, this looks just like a drawing tool. Apart from the semantics of the boxes and their possible relationships, the application itself does not seem to be embedded with the knowledge on the topic that would allow me (the non-expert) to quickly discover the range of vulnerabilities in my current system design. [...]

I think it needs to seek a better problem-solution fit, but if you get it right it would be a really helpful as a tool [sic].

—*Expert 4*

Section 14.1 discusses the experts' feedback as well as the students' comments. The latter ones go along with the answers provided by the experts: S4, S5, S7, and S9 explicitly deemed the privacy model as "comprehensive," while all remaining students used similar expressions. Furthermore, S5 added that the model would need constant updating, since "new threats may appear in the future." S1, S2, S4–S8, S10, and S12 explicitly said that the threat model may actually support designers of public display systems, for example:

The threat model can assist designers by providing a complete set of threats, weaknesses and counter measures [sic]. This way designers can pay attention to circumstances they might not have thought of. —*Student 1*

Designers are able to check if there are any threats for their systems they did not think of in first place. It is also possible to lookup [sic] counter measures [sic]. —*Student 5*

Designers could test their design in a systematic way by using the model as a reference. —*Student 12*

The remaining three students did not address designers in particular, but pointed to more general characteristics of the prototype:

It makes people think about the issues in a structured way. —*Student 3*

One can draw up various scenarios and possible threats beforehand. Knowing as much as possible scenarios is essential for developing a system that covers most of the privacy leaks. —*Student 9*

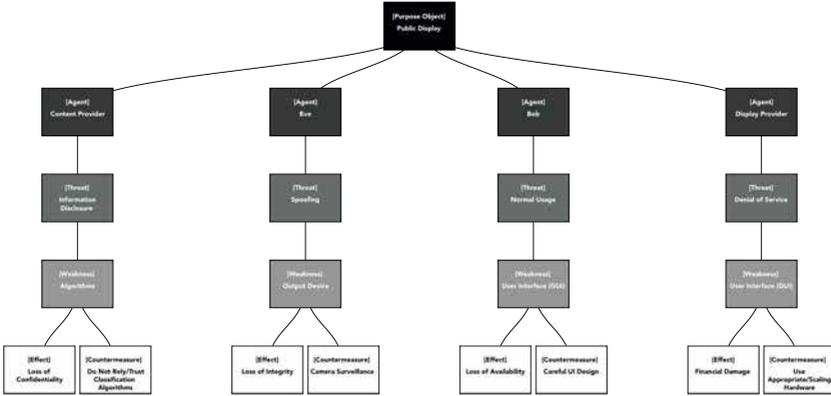
The model is a good way for visualizing a brainstorming. It helps to see threats, weaknesses etc. —*Student 11*

Finally, some students also provided general comments on the privacy threat model or its corresponding prototypical implementation, for example:

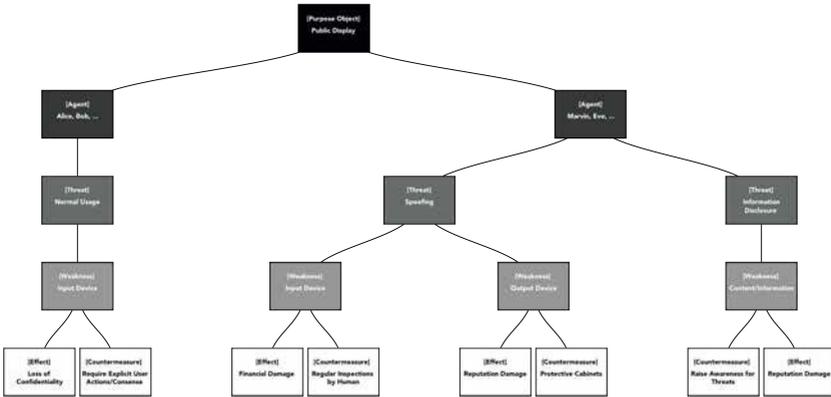
It creates a coherent framework to analyze a system and reflect on potential privacy impacts. —*Student 7*

At the first glance, the six threat models created by the students exhibit two patterns, see Figure 12.1: breadth-first and depth-first. Figures 12.1a and 12.1b are examples for the first one, while Figures 12.1c and 12.1d are examples for the latter one. Disregarding the negligible right trunk, Figure 12.1e also resembles a depth-first approach. A closer look reveals, that the breadth-first approaches include both, malicious and benign users, i.e., Alice or Bob as well as Marvin or Eve. In contrast, the depth-first approaches start their privacy threat analysis with malicious users first. The threat model depicted in Figure 12.1f appears to be an exception, as it only features neutral threat agents. Section 14.1 discusses this finding of the breadth-first and the depth-first approach in more detail.

Another finding is that four threat models identify spoofing as a major privacy threat, and five of the six threat models name information disclosure as a potential privacy threat. Both findings correspond with the results presented in Section 10.1. Though the comments indicated that the threat model was perceived as comprehensive, the threat model shown in Figure 12.1f introduces a custom countermeasure, i.e., “check code etc.” It is used to mitigate the negative effect of loss of confidentiality caused by human failure (presumably programming errors). Eventually, four of the six threat models introduce a second purpose object, i.e., a countermeasure that may be subject to a specific privacy threat in turn. However, only three of them present a complete second iteration. Figure 12.1f stops right after defining the second purpose object and after defining the specific threat.



(a)



(b)

Figure 12.1.: The students' privacy threat models. Created with the IPDPTM and edited afterwards to improve readability.

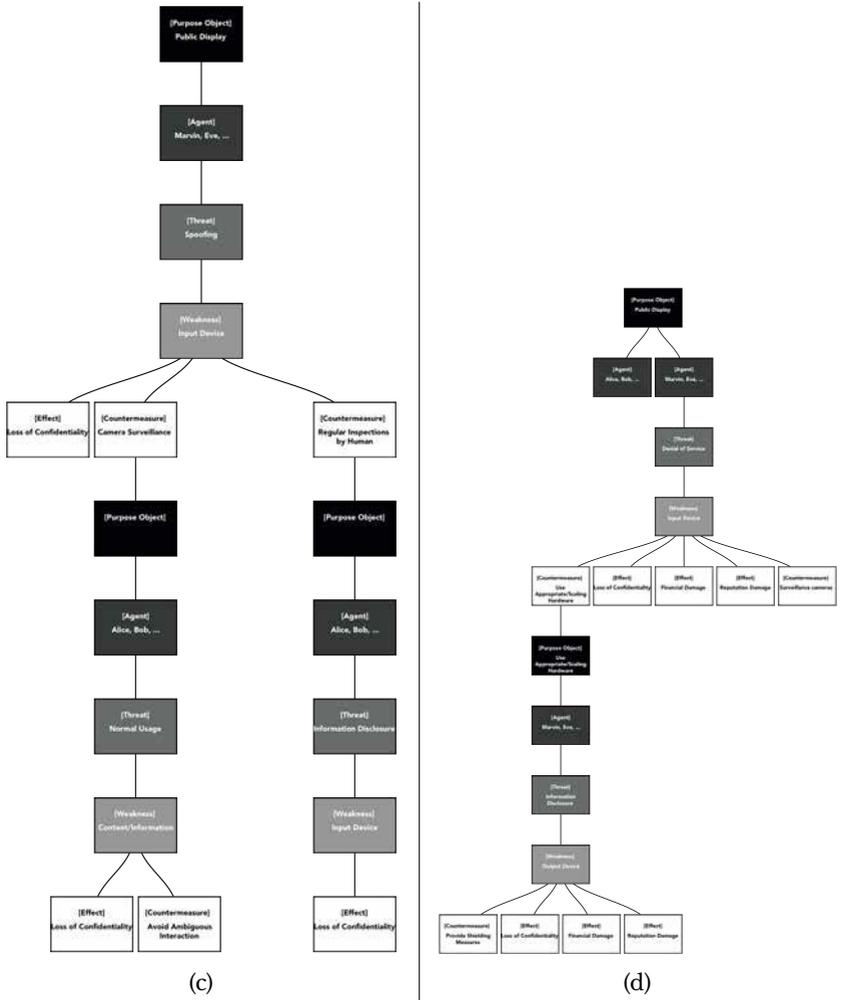
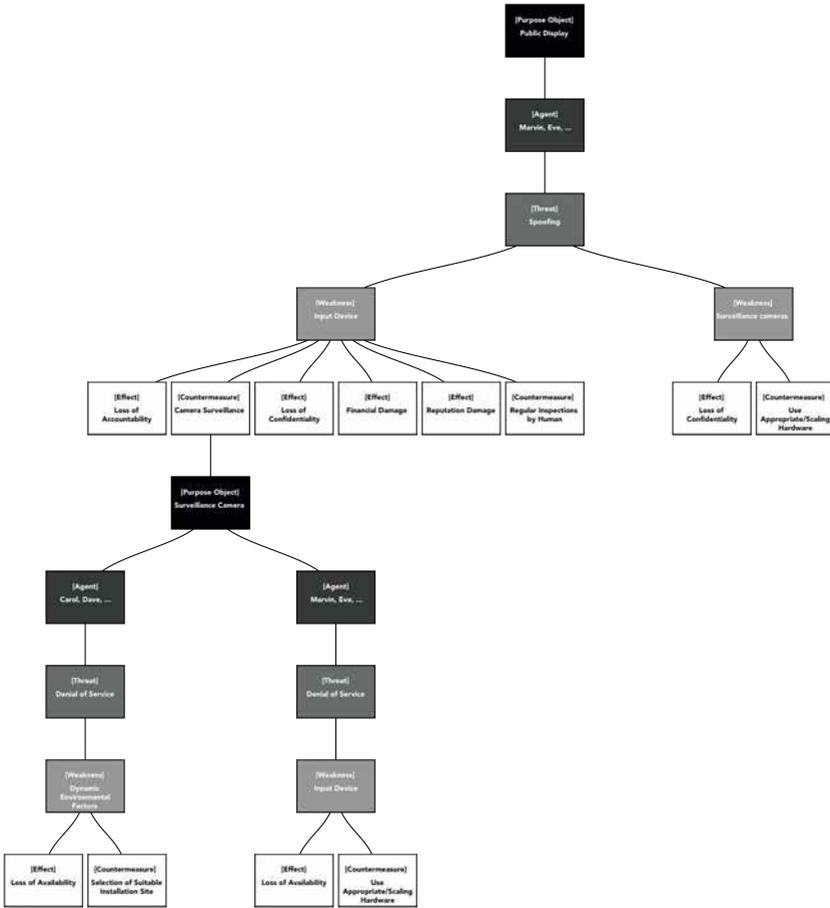
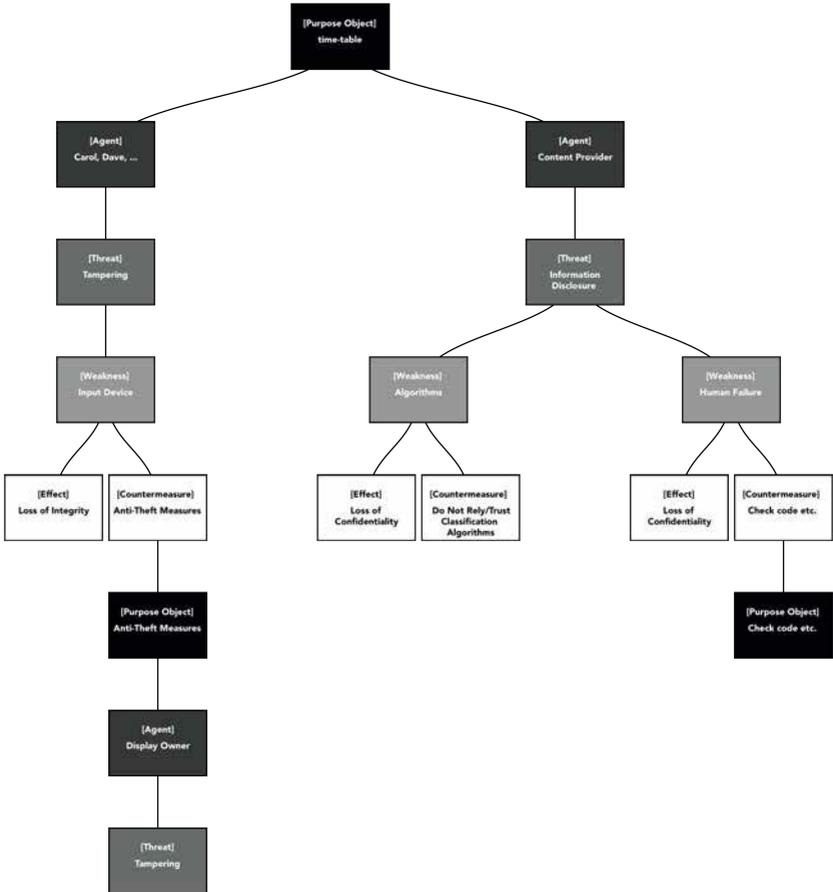


Figure 12.1.: The students' privacy threat models (continued).



(e)

Figure 12.1.: The students' privacy threat models (continued).



(f)

Figure 12.1.: The students' privacy threat models (continued).

Table 12.1.: Results of the NASA TLX questionnaire (M: mental demands, P: physical demands, T: temporal demands, PF: performance, E: efforts, F: frustration).

M	P	T	PF	E	F
1.33	-4.67	-1.17	-1.33	-1.17	-0.25

12.2. Countermeasures

Three novel countermeasures were proposed in Section 10.2, i.e., visual multiplexing, visual highlighting, and visual interaction. Section 11.2 presented three prototypical implementations based on these approaches. These prototypes were used to evaluate the approaches as follows: Subsection 12.2.1 presents the results of a user study that was carried out to assess each visual multiplexing method with regard to perceived mental workload, suitable contents, and technical feasibility. Subsection 12.2.2 focuses on the results of a user study conducted to analyze the efficiency, effectiveness, and robustness of visual highlighting. Eventually, Subsection 12.2.3 reports on the results of a study carried out to evaluate visual interaction with respect to maximum operating distance, hardware impacts, and multi-user potential.

12.2.1. Visual Multiplexing

Throughout the development of the prototype, all multiplexing methods were tested with different content types. These tests included maps, photos, symbols, words, and short sentences. Figure 12.2 shows some examples. During these tests, it became apparent that some multiplexing methods worked well for certain content types, but not

for others. The user study presented in this subsection was thus designed to gain further knowledge on this issue. The aim was to (i) compare the raw performances of FDM, CDM, and TDM in terms of their suitability for different content types and to (ii) deduce the characteristics of each method based on these results. To focus on the raw performance only, the study was lab-based and used the entire public display for the multiplexed content, rather than only parts of it.

In order to produce results that provide answers towards RQ2 (“What are countermeasures to those privacy threats?”) and its subquestion “Do countermeasures impact the general public display usage?”, the study aimed at the questions and hypotheses listed in Table 12.2. Furthermore, general feedback on user satisfaction and usability, as well as general comments on each multiplexing approach were recorded.

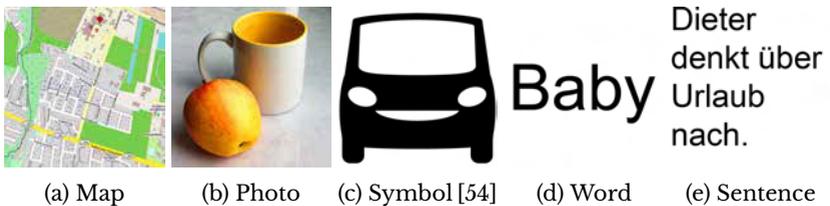


Figure 12.2.: Examples of content types used as input images in the study about visual multiplexing.

Apparatus and Material

The study was conducted in a laboratory environment, having a 52” LCD screen simulate a public display. The multiplexed content was prepared as a slide show. The mobile device was an Apple iPad 3, running iOS 5.1.1 and version 2.2 of the prototype. The iPad 3 was chosen over other devices as it provided superior processing power at the

Table 12.2.: Questions and hypotheses used in the user study on visual multiplexing.

ID	Description
Q1	How well can users make use of visual multiplexing on public displays?
Q2	Which multiplexing method is most suited for a specific type of content?
H6	Symbols work well with all multiplexing methods (due to their simple design and messages).
H7	Maps will be handled poorly by all multiplexing methods (due to their high information density).
H8	TDM is the most accurate multiplexing method.

time the study was conducted. The study was based on a paper-based questionnaire, that has been divided into five parts: an introduction with an initial demographic questionnaire and three further blocks in identical order, one for each multiplexing method. Every block was subdivided into five sections, each featuring one content type, again in identical order. A NASA TLX questionnaire was put at the end of each section. The fifth part of the questionnaire provided space for any additional written feedback the participants wanted to give.

The five content types depicted in Figure 12.2 were chosen based on Schaeffler's [195] work, as they were considered as good examples for the kind of content commonly shown on public displays. The study material was created in the following way: All content types were prepared as bitmaps of 500 x 500 pixels in size. Thirty roadmap tiles were taken from OpenStreetMap¹. The tiles show maps of randomly picked places in Germany at zoom level 15 (see Figure 12.2a). Flickr²

¹<http://www.openstreetmap.de>, accessed: January 22, 2013.

²<http://flickr.com>, accessed: January 22, 2013.

served as a source for randomly picked photos. Thirty of the “most recent photos and videos” were converted into tiles by clipping out their centers (see Figure 12.2b). The symbols were taken from The Noun Project³, a website that hosts black and white symbols depicting objects. Thirty symbols were randomly chosen and converted into bitmaps with a white background (see Figure 12.2c). Thirty short German words were randomly picked from a dictionary and thirty short German sentences were built, each three to six words in length. The words and sentences were rendered in a black 100 pt Helvetica font on a white background (see Figures 12.2d and 12.2e).

Participants

In total, 21 people (16 male, 5 female) participated in the study. They were recruited via e-mail, mailing lists, bulletin boards, and social networks. Their age ranged between 20 and 40 years. Participants included students, office workers, teachers, and PR managers.

Procedure

The study was conducted in German, with two experimenters present in the room at all times. One experimenter lead the study, while the other took notes on comments and actions. The study used a within-subject design, thus every participant completed all 15 sections (3 multiplexing methods \times 5 content types). Participants completed the study in approximately 15 minutes.

After the participants were welcomed, the experimenters briefed them about the study and asked them to fill in the initial questionnaire gathering demographic information. During the main part of study, the

³<http://thenounproject.com>, accessed: January 22, 2013.

participants were told to stand approximately one meter in front of the public display, while facing it frontally. As suggested by Boring et al. [35], placing participants about 1.0–1.5 meters away from the display provides a comparable viewing impression as if standing 3.0–6.0 meters away from a typically larger public display.

In the main part of the study, participants had to carry out a number of matching tasks grouped into three blocks, i.e., one for each of the multiplexing methods. One experimenter set up the mobile device in preparation for each block. Each block consisted of five sections—one for each content type—and for each of these sections, participants had to first perform a matching task and then to assess their workload using the NASA TLX questionnaire.

Participants were asked to hold the mobile device towards the public display and to look at the demultiplexed content. They were then asked to match the shown content to one of six choices on the answer sheet by marking the corresponding option. After the matching task, the participants filled in a NASA TLX form. This procedure was repeated for all content types, and for all multiplexing methods. After completing the main part, the experimenters asked participants for any further comments they might want to provide, debriefed them, and paid them a small sum of money for participating.

Results

The demographic questionnaire revealed that all participants owned a mobile phone and had some experience with touch-enabled devices. The participants' age did not have a significant effect on the number of correct answers, as indicated by a oneway ANOVA ($F(7, 315) = 2.04, p > 0.161$). There was also no correlation between the participants' experience with mobile phones and the number of correctly recognized

images ($F(4, 315) = 2.40, p > 0.337$). Overall, the participants were able to correctly identify 84.44% of all images.

Content Types. Figure 12.3 shows the correlation between the number of correct answers and the content types for each visual multiplexing method as well as in total. Overall, only 63.49% of all maps were identified correctly. Most participants had difficulties to distinguish the maps shown with FDM (38.10% correct answers) and CDM (57.14% correct answers). TDM performed significantly better with 95.24% of correct answers. In contrast to maps, users could recognize 98.41% of all symbols in total. 52.38% of the words and 57.14% of the sentences were recognized by the participants using FDM. 76.19% of the photos were correctly identified if FDM was used, 95.24% if CDM was used, and 100.00% if TDM was used. A oneway ANOVA supports a significant effect of the content type on the number of correctly identified images ($F(4, 310) = 2.40, p < 0.001$).

Multiplexing Methods. Figure 12.4 shows how many images were identified correctly for each multiplexing method. It also breaks this number down for every FDM information channel, since the study randomly changed the FDM information channel for each participant, i.e., eight participants used the red, seven the green, and six the blue information channel. The results are in favor of the red information channel: 53.73% of the images were identified correctly, whereas green and blue scored less, 20.90% respectively 25.37%.

To investigate this phenomenon further, the captured camera image was analyzed using a typical TV test image (an EBU 75/75% test card⁴). The test card was shown on the public screen and photographed with

⁴http://en.wikipedia.org/wiki/Test_card, accessed: May 20, 2015.

the camera of the mobile device. The captured RGB values for the red, green, and blue test bars were then analyzed. The camera actually produced the most accurate results for the red test bar: The expected triple was (255, 0, 0) and the captured triple was (242, 31, 15). Looking at the green test bar, the captured triples were significantly worse: (0, 255, 0) was expected, but the camera returned (0, 247, 153) instead. The results for the blue test bar were better: (0, 0, 255) was expected, but (0, 19, 255) was captured. A oneway ANOVA confirmed the effect of the chosen FDM information channel on the number of correctly recognized images ($F(2, 102) = 3.09, p < 0.001$).

Using FDM, only 63.81% of all answers were correct. Compared to FDM, the results for CDM are better for photos, symbols, words, and sentences. Using TDM, 99.05% of all maps, photos, symbols, words, and sentences were recognized correctly. A oneway ANOVA affirms a significant effect of the multiplexing method on the number of correctly identified images ($F(2, 312) = 3.03, p < 0.001$).

Workload. The NASA TLX results relate to values on a scale from 1 (“very low”/“perfect”) to 20 (“very high”/“failure”). The participants were told to interpret the question about temporal demands in such a way that higher values are better (“the system is fast”) than lower values (“the system is slow”). To harmonize this with the remaining NASA TLX scores, the results for temporal demands were inverted. Figure 12.5 shows the median values for all NASA TLX scores by content type. Figure 12.6 visualizes the median values for all NASA TLX scores by multiplexing method.

An analysis of the NASA TLX results and the correct answers given shows that in 79.37% of all cases participants that rated their performance between 1 and 10 in the NASA TLX, had indeed given the correct answer. Only 9.84% believed to be right, when in fact, they were

wrong. When looking at the NASA TLX results by visual multiplexing method and content type (not depicted), the CDM results are particularly striking, as the performance for maps is remarkably bad: Participants claimed high demands for almost all categories (mental demands: 11, temporal demands: 10, performance: 13, effort: 12, and frustration: 9). In contrast, CDM was favored for symbols (mental demands: 3, temporal demands: 4, performance: 1, effort: 3, and frustration: 2) and words (mental demands: 3, temporal demands: 4, performance: 1, effort: 3, and frustration: 2). This concludes the evaluation of visual multiplexing as to presenting the results obtained from the user study. Subsection 14.2.1 is going to discuss the presented results.

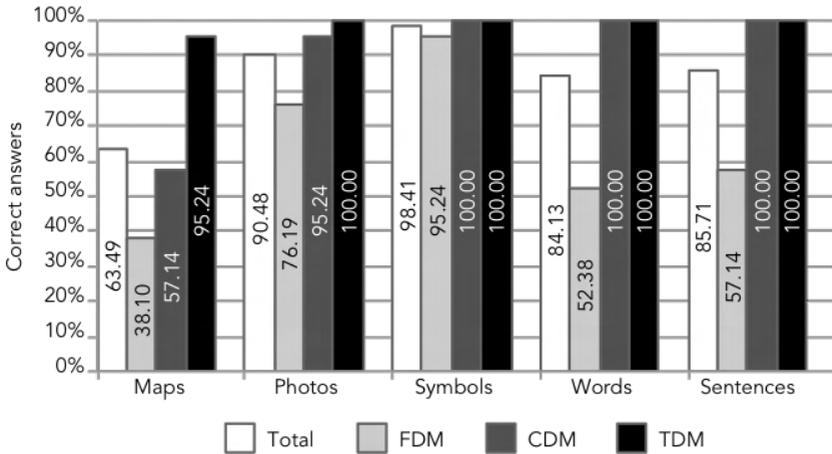


Figure 12.3.: Correct answers by content type.

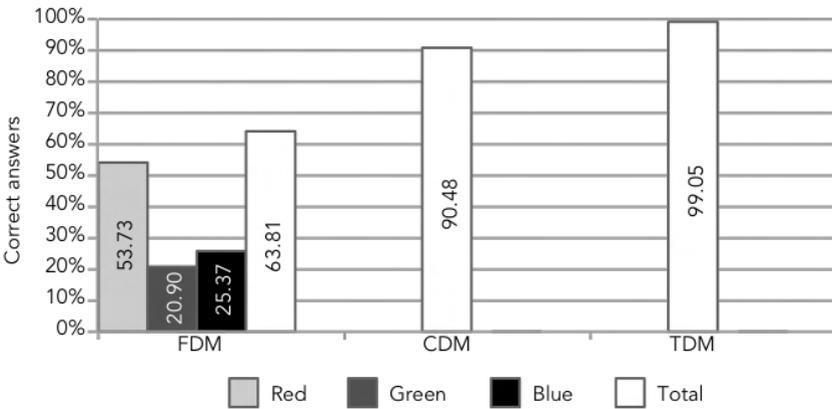


Figure 12.4.: Correct answers by method; FDM shows percentage of correct answers per color channel.

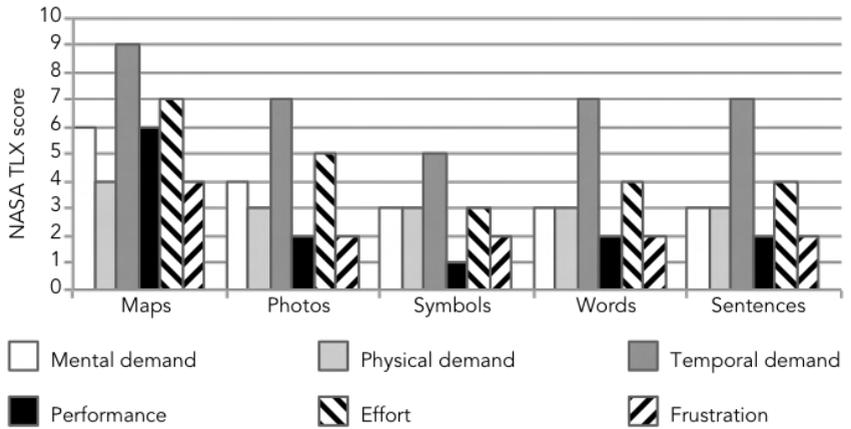


Figure 12.5.: NASA TLX medians by content type.

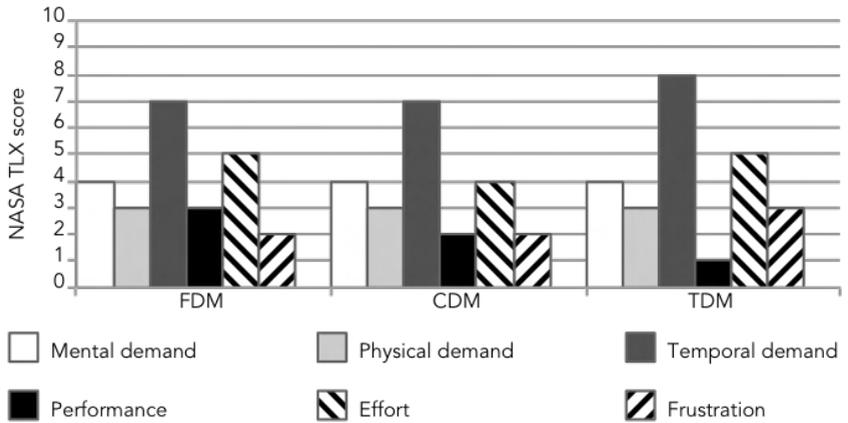


Figure 12.6.: NASA TLX medians by multiplexing method.

12.2.2. Visual Highlighting

Subsection 10.2.2 proposed a set of comparison criteria, that can be used to quantify the characteristics of different approaches towards highlighting on public displays. Table 10.14 provides an overview of all criteria and the corresponding scales. The remainder of this subsection compares the approach presented in this thesis to other existing approaches. Afterwards, this subsection presents the results of a user study, which was conducted to assess the raw performance of visual highlighting with regard to efficiency, effectiveness, and robustness.

For a quick navigation within the provided data and to provide a systematic and equal comparison scheme to every approach, the findings are presented in Table 12.3 (quality), Table 12.4 (quantity), and Table 12.5 (reliability and robustness). The findings are also visualized as a spider diagram in Figure 12.7. For the sake of readability, this diagram only shows the comparison criteria, whose values differ the most. One important result that can be drawn from the spider diagram is that the area spanned by the prototype presented in Subsection 11.2.2 exceeds the areas spanned by all other projects. As larger values correlate to better characteristics on the visualized axes, this indicates that the approach to visual highlighting on public displays as presented in this thesis surpasses all existing approaches with regard to the given comparison criteria. The comparison table only contains data for comparison criteria, whose values can be derived from the original publications mentioned in the related work, see pp. 222.

In addition to the comparison presented above, a user study was carried to assess the performance of the Multipleye prototype with regard to visual highlighting. The focus of the study was in particular on clarity, granularity, delay, readability, and correctness. However, directly measuring these criteria appears challenging, which is why

Table 12.3.: Comparison of highlighting approaches with regard to quality (see Table 10.14 for units).

Project	Clarity	Granularity	Duration	Delay	Readability
CrossFlow [43]	Unknown	3, an arbitrary direction	2, pre-defined by time taken to highlight one information multiplied by amount of distinct information	1	1
CrossBoard [43]	Unknown	3, ranges from block to pixel	2, pre-defined by time taken to highlight one information multiplied by amount of distinct information	1	2
Rotating Compass [189]	Unknown	1, one direction out of 4 pre-defined, finer and distinct set of directions	2, pre-defined by time taken to highlight one direction multiplied by number of distinct directions	1	1
Interactive Ambient Public Displays [231]	Unknown	2, ranges from blocks to words according to user's positioning	Not applicable	1	2
Screen Codes [58]	Unknown	Not applicable	Not applicable	1	1
SnapAndGrab [136, 137]	Unknown	Not applicable	3	10 (estimation)	2
Multipleye	Unknown	6, pixel	3	5	3

Table 12.4.: Comparison of highlighting approaches with regard to quantity (see Table 10.14 for units).

Project	Interference	Concurrent highlighting	Concurrent users	Time density
CrossFlow [43]	3	1	Infinite	1 Hz, estimation. Depends on minimum time the user needs to sense and process the crossmodal cue.
CrossBoard [43]	3	Half of the total amount of information	Infinite	1 Hz, estimation. Depends on minimum time the user needs to sense and process the crossmodal cue.
Rotating Compass [189]	3	1	Infinite	1 Hz, estimation. Depends on minimum time the user needs to sense and process the crossmodal cue.
Interactive Ambient Public Displays [231]	1	3, depends on the number of concurrent users. Paper implies 2-3.	3, paper does not state explicitly	Not applicable
Screen Codes [58]	1	Not applicable	Infinite	Paper does not provide information about this
Snap and Grab [186, 187]	3	Not applicable	Infinite	Paper does not provide information about this
Multipleye	3	4, depends on the size of the QR/AR tag	Infinite	Depends on the size of the used QR/AR tag. In average, one tag can hold 2-4 parallel information. If multiple QR/AR tags are necessary, this value is defined by the time the smartphone requires to scan each tag.

Table 12.5.: Comparison of highlighting approaches with regard to reliability and robustness (see Table 10.14 for units).

Project	Availability	Correctness	Environmental influences
CrossFlow [43]	Ready after each cycle	Neutral	May cover the crossmodal cue or may be erroneously sensed as the crossmodal cue. Users may be unable to sense the crossmodal cue, e.g. due to numbness.
CrossBoard [43]	Ready after each cycle	Neutral	May cover the crossmodal cue or may be erroneously sensed as the crossmodal cue. Users may be unable to sense the crossmodal cue, e.g. due to numbness.
Rotating Compass [189]	Ready after each cycle	Neutral	May cover the crossmodal cue or may be erroneously sensed as the crossmodal cue. Users may be unable to sense the crossmodal cue, e.g. due to numbness.
Interactive Ambient Public Displays [231]	Instantly if number of users blow maximum	Unlikely, as based on personal profiles	Display may relay out the screen when users appear or disappear
Screen Codes [98]	Always, personal mobile device requires time to scan and process code.	Not applicable	Recorded image may be distorted, e.g., by tearing.
Snap and Grab [186, 187]	Always, picture taking, processing, and transmission takes time.	Unlikely, as system returns information selected by the user	Users may not be able to take a picture of the desired display area. Reflections may render the taken pictures undetectable.
Multipleye	Always, personal mobile device requires time to scan the QR/AR code tag.	Unlikely, as highlighting is tightly bound to the actual information	Code tag may be un-scannable due to physical constraints, e.g. distance to the display.

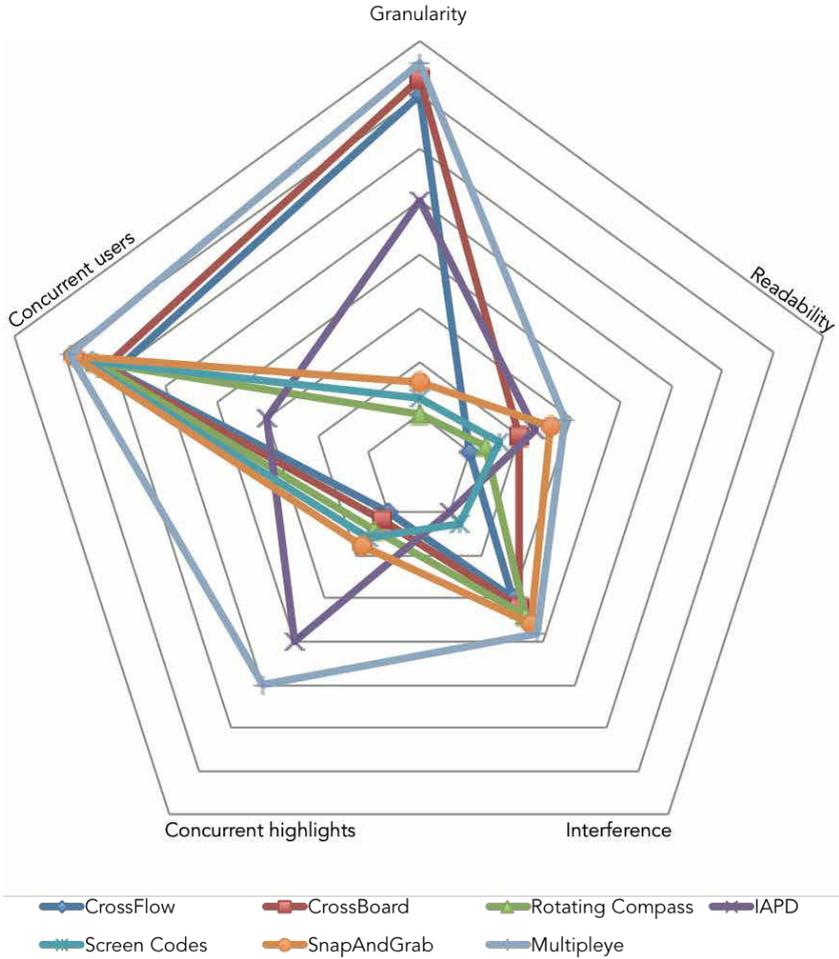


Figure 12.7.: Spider diagram showing selected results of the comparison shown in Table 12.3, Table 12.4, and Table 12.5. Lower values are shown towards the center. For a detailed description refer to the corresponding table.

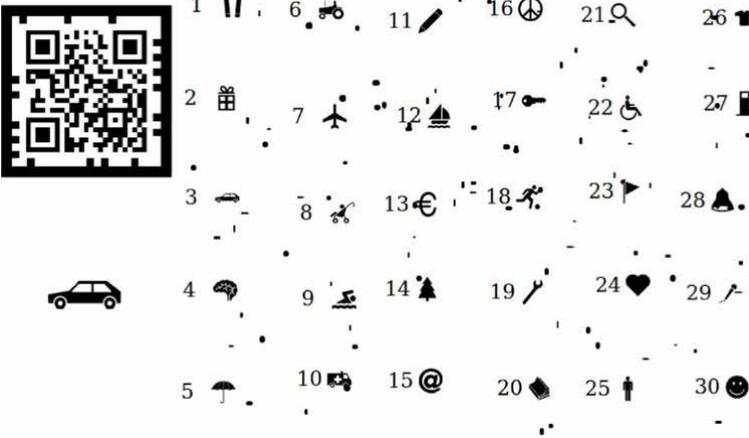
surrogate values were used. Some of the surrogate values are aggregations of the aforementioned criteria: The clarity and readability of the approach could be consolidated in the term effectiveness; granularity and delay could be summarized as efficiency; the correctness of the system correlates to the robustness. Eventually, these surrogate values were measured in the user study to find answers to the questions listed in Table 12.6. The same table also introduces the formulas used to calculate the values for efficiency, effectiveness, efforts, and performance (all values without units).

Achieved objective corresponds to the number of correctly identified pieces of information, while *defined objective* refers to the maximum number of pieces of information that had to be identified. The number of identified pieces of information can be measured objectively, as can identification times. The subjective temporal demands were thus left out of the NASA TLX questionnaire. Mental and physical demands, however, can only be assessed via the participants' subjective gradings on such a questionnaire. There are 20 values to choose from on a NASA TLX scale, so each step equals $0.05 = 5\%$ (see definition of efforts in Table 12.6), that add up to 100% from far left to right.

The robustness of the system was assessed by adding visual noise to the contents shown on the public display. This noise was designed as black, unevenly shaped flecks, spread throughout the content in a random manner. Even though QR codes are designed to be resistant to such noise, the flecks did not appear within the QR codes. The rationale for this is that the study aims to measure the robustness of visual highlighting on public displays rather than the robustness of QR codes. In addition to the flecks, the icons were also randomly moved out of the grid. Figure 12.8a shows normal example used in the study, while Figure 12.8b shows the described visual noise.



(a) Without visual noise



(b) With visual noise

Figure 12.8.: Grid of icons, which was used in the user study about the performance of the prototypical implementation of visual highlighting.

Table 12.6.: Research questions and calculations used in the user study about the performance of the prototypical implementation of visual highlighting (VH).

ID	Research question
Q3	Does VH have a positive effect on efficiency when retrieving information on public displays? $\text{efficiency} = \frac{\text{performance}}{\text{efforts}} := \text{EF1}$ $\text{performance} = \frac{\text{fastest identification time of all participants}}{\text{mean identification time of one participant}}$ $\text{effort} = 0.05 * \frac{\text{mental demands} + \text{physical demands} + \text{endeavors}}{3}$
Q4	Does VH have a positive effect on effectiveness when retrieving information on public displays? $\text{effectiveness} = \frac{\text{achieved objective}}{\text{defined objective}} := \text{EF2}$
Q5	Does VH have a positive effect on robustness when retrieving information on public displays?

Apparatus and Material

A 52" LCD screen running at 1920 x 1080 pixels simulated a public display. As Boring et al. [35] suggest, placing the participant about 1.0 to 1.5 meters away from this screen provides a comparable viewing impression as if the participant stood 3.0 to 6.0 meters away from a common public display. A Mac mini showed a slideshow with the prepared content. Additionally, the Mac mini could be controlled by the experimenter and the participant as explained below. The mobile device was simulated by an iPad 2 running iOS 5.1. The tablet was chosen over a phone-sized device to accommodate the additional timekeeping controls used within the study (see below).

The software on the personal mobile device was based on the publicly available prototype, see Subsection 11.2.2. To allow for precise timekeeping, a start and a stop button was added to the user interface in order to control an internal stopwatch. Furthermore, the personal mobile device was connected to the public display in order to control the content as follows: At first, the public display shows a black screen only. As soon as the participant hits the start button, the stopwatch starts and the public display reveals the actual content of the study. When the participant hits the stop button, the stopwatch shows the measured time and the public display turns black again. This prevents the participants from searching after they stopped the stopwatch.

The study itself was designed as a questionnaire, which consisted of two blocks: The first block aimed at measuring the efficiency and the second block was designed to measure the effectiveness. Each block was further subdivided into four sections. The first two sections did not use the personal mobile device, whereas the second two sections did. To test for the robustness of the system, every other section used visual noise. Thus, the conditions, i.e., visual highlighting and visual noise, were counterbalanced. Table 12.7 gives an overview of the structure of the questionnaire and also provides an abbreviation for each section for later reference.

Sections 1–4 contained eight tasks each, sections 5–8 contained one task each. There was an example for every task at the beginning of each section. To estimate the participants' perceived efforts, a NASA TLX sheet was included at the end of each section. Thus, a questionnaire consisted of 36 tasks and 8 NASA TLX sheets. Additionally, demographic questions, for example, about the participants' age, gender, or experience with tablet computers, were asked in a pre-questionnaire. To speed up the data gathering and subsequent evaluation, the entire questionnaire was designed as a computer based form,

that was either filled in by the experimenter or, in case of a NASA TLX sheet, directly by the participants.

The experimenter prepared 87 slides, each comparable to the ones shown in Figure 12.8, except for a different set of icons. In all, 35 icons were taken from a website⁵, each visualizing a commonly-known object, e.g., a pencil, a heart, or a key. All icons were unified in size (40 x 40 pixels) and randomly placed on the slides. Each slide also showed a code tag in the upper left corner and a grid of varying icon subsets, namely 30 icons, in the remaining space. An adjacent number, ranging from 1 to 30, identified every icon. Pre-tests revealed that locating a particular icon amongst this number of icons represents a reasonable challenge to the participants. The icon that the participants had to locate was randomly defined for each slide. Icons were chosen over other content types, as icons resemble many public display scenarios, e.g., pedestrian navigation. Furthermore, it is beneficial for the study, as the distinctive shape of icons is clearly recognizable, even for impaired, e.g., color blind, participants. Though there may be many more types of visual noise that could appear on a public display, the noise applied to the slides in the study should already serve as a valid indicator for the robustness of the system.

Table 12.7.: Overview of the sections used in the user study.

Section	Abbreviation	Highlighting	Noise
1, 5	\overline{HN}	no	no
2, 6	\overline{HN}	no	yes
3, 7	$H\overline{N}$	yes	no
4, 8	HN	yes	yes

⁵<http://thenounproject.com>, accessed: May 27, 2015.

Participants

The experimenter recruited 24 participants by e-mail, bulletin boards, and social networks. 19 of them were male, 5 female. Their age ranged between 20 and 40 years with an average of 27.5 years ($s = 5.56$). The participants were students, office workers, teachers, and PR managers.

Procedure

One experimenter conducted the study indoors in a lab environment. Due to the within-subject design, every participant completed all 36 tasks and filled in all 8 NASA TLX sheets. On average, a participant completed the study in 20 minutes. At the beginning of each session, the participants were briefed and asked to answer the questions in the pre-questionnaire. The experimenter prepared the tablet device as required for each section, so that the device could be used either for timekeeping purposes or as the personal mobile device used for the visual highlighting. Afterwards, the experimenter explained each following task to the participant and asked for further questions.

The procedure for sections 1 ($\overline{H\bar{N}}$) and 2 ($\overline{H}N$) was as follows: The participants were shown the first of 16 icons (8 per section) that they were to locate. Once they memorized the icon, they could press the start button on the tablet device to start the stopwatch and uncover the public display. In these two sections, the participants only used their bare eyes to complete the task and did not use visual highlighting. As soon as the participants spotted the location of the icon, they pressed the stop button in order to stop the stopwatch and cover the public display again. The tablet device then showed the measured time in seconds with an accuracy of a hundredth of a second. The participants reported the number they identified next to the searched icon to the

experimenter. If the number was correct, the experimenter recorded the time according to the tablet device and prepared the system for the next task. If the number was wrong, the participants had to redo the task again, but the stopwatch was not reset. This way, only correct answers were collected. Once the participants finished a section, they filled in the NASA TLX sheet.

The procedure for sections 3 ($H\bar{N}$) and 4 (HN) was comparable to the procedure of the previous two sections, except for that the participants now used the tablet device as their personal mobile device for visual highlighting to complete the task. Therefore, the participants held the tablet in such a way that they could see the content of the public display on the screen of the tablet. As soon as the visual highlighting appeared, the participants memorized the corresponding number and pressed the stop button. The study proceeded as described above.

The procedure for sections 5 through 8 was as follows: The experimenter set a timer for 20 seconds. Once these 20 seconds elapsed, an alarm would sound, telling the experimenter and the participants that the task has ended. As soon as the participants started the timer, the content of the public display was revealed. As soon as the participants located the requested icon, they called out the according number. If that number was correct, the experimenter moved on to the next icon. If the number was wrong, the participants had to keep searching. At the end of the task, i.e., after 20 seconds, the number of correctly identified icons was recorded. Once the participants finished each section, they filled in the NASA TLX sheet. The differences between sections 5 through 8 is in the varying use of visual highlighting and the applied visual noise, see Table 12.7 for reference. At the end, the participants were asked for any additional or concluding comments and were debriefed. The experimenter compensated all participants for their time by paying them a small amount of money.

Results

The results of the pre-questionnaire showed that only 8.33% of the participants owned a tablet device. 29.17% claimed to be unfamiliar with such a device, 50% of the participants said that they had at least some experience. In contrast to the acquaintance with a tablet device, 72% of the participants stated to be quite familiar with touchscreens.

Efficiency. An analysis of the recorded measurements reveals an apparent learning effect for the first three icons. Thus, these values were suppressed in the further evaluation. The corrected mean values calculate to: 2.38 seconds for section $\overline{H\overline{N}}$ and 1.74 seconds for section $H\overline{N}$, see Table 12.8. A oneway ANOVA confirmed a medium effect of the section on the required time ($F(3, 624) = 2.62, p < 0.0001, r = 0.43$).

Effectiveness. In average, the participants located 4.54 symbols in section $\overline{H\overline{N}}$ and 4.88 symbols in section $H\overline{N}$, see Table 12.9. Thus, the applied visual noise roughly reduces the speed by 1 symbol per minute (0.333 symbols in 20 seconds). Using visual highlighting, participants were able to locate about twice as many symbols: 10.71 symbols in section $H\overline{N}$ and 11.79 symbols in section HN . A oneway ANOVA confirmed a large effect of the section on the number of located symbols ($F(3, 96) = 2.71, p < 0.0001, r = 0.94$). According to a Tukey-HSD post-hoc analysis, the visual noise has a significant influence on the results of the study between sections $\overline{H\overline{N}}$ and $H\overline{N}$ ($MD = -0.83, p = 0.000$). However, there is no significant influence on the results between sections $H\overline{N}$ and HN ($MD = -0.33, p = 0.799$).

Workload. Looking at the sections about efficiency, the NASA TLX records show comparable values for mental demands: 3.71 ($\overline{H\overline{N}}$), 5.00

(\overline{HN}), 3.08 ($H\overline{N}$), and 3.29 (HN), see Table 12.10. Performance, efforts, and frustration also only show minor differences. The recorded values about effectiveness, however, reveal a noticeable difference in mental demands: Visual highlighting reduces the mental demands to about half as much as in sections without that approach, i.e., 8.75 (\overline{HN}) and 8.50 ($H\overline{N}$) in contrast to 3.13 ($H\overline{N}$) and 3.42 (HN). The results for efforts and frustration behave the other way around: 10.83 (\overline{HN}) and 10.04 ($H\overline{N}$) in contrast to 3.96 ($H\overline{N}$) and 3.83 (HN) for efforts, 4.96 (\overline{HN}) and 4.75 ($H\overline{N}$) in contrast to 1.96 ($H\overline{N}$) and 2.00 (HN) for frustration. This concludes the evaluation of visual highlighting as to presenting the results obtained from the study. Subsection 14.2.2 is going to discuss the presented results.

Table 12.8.: Results of the efficiency block (unit: seconds, EF1: efficiency, see Table 12.6).

Section	\bar{x}	M	s	min	max	EF1
\overline{HN}	2.38	1.90	1.44	1.05	11.62	2.85
$H\overline{N}$	3.22	2.64	2.20	0.87	17.05	1.91
$H\overline{N}$	1.74	1.55	0.66	0.93	4.72	3.55
HN	1.69	1.58	0.45	0.96	3.76	3.80

Table 12.9.: Results of the effectiveness block (unit: number of icons, EF2: effectiveness, see Table 12.6).

Section	\bar{x}	M	s	min	max	EF2
\overline{HN}	4.54	5	1.69	7	1	0.65
$H\overline{N}$	4.88	5	1.04	6	3	0.70
$H\overline{N}$	10.71	11	1.33	13	8	1.53
HN	11.79	12	0.83	13	10	1.68

Table 12.10.: Results of the NASA TLX questionnaire (M: mental demands, P: physical demands, PF: performance, E: efforts, F: frustration).

Section	M	P	PF	E	F
\overline{HN}	3.71	2.42	14.63	7.63	3.29
\overline{HN}	5.00	3.04	14.63	7.17	3.04
$H\overline{N}$	3.08	6.00	15.04	6.04	3.25
HN	3.29	4.96	15.38	6.29	2.54
\overline{HN}	8.75	2.58	9.54	10.83	4.96
\overline{HN}	8.50	2.96	10.21	10.04	4.75
$H\overline{N}$	3.13	3.67	15.75	3.96	1.96
HN	3.42	3.41	16.08	3.83	2.00

12.2.3. Visual Interaction

To assess the general feasibility and baseline performance of Lichtblick, which constitutes the prototypical implementation of the approach to visual interaction as presented in Subsection 11.2.3, a lab-based study was conducted. Table 12.11 lists the characteristics that the study focused on in particular.

Table 12.11.: System characteristics addressed by the study.

ID	Description
Q6	Maximum operating distance
Q7	Hardware impacts, e.g., smartphone brand
Q8	Multi-user potential

Apparatus and Material

The system was tested in a lab environment. The first component, i.e., the public display, was simulated by a generic TFT. The specific characteristics of the display have little impact on the performance of the system, as it is merely used to display QR codes besides regular content. The camera used to capture the light signals was a Sony EyeToy, with a resolution of 640 x 480 pixels and a framerate of 30 fps. This camera was chosen since it is an affordable consumer product with sufficient characteristics, that supports the idea of a lightweight and simple deployment. The test environment was a standard office with overhead lighting (fluorescent lamps) and shut blinds to ensure controlled conditions throughout the test. The second component, i.e., the smartphone application, was tested on the following, opportunistically chosen devices: Samsung Nexus S (D1), Samsung Galaxy Beam

(D2), HTC Legend (D3), HTC Desire (D4), and HTC One S (D5). The devices were running Android versions from 2.2 up to 4.1.

Procedure

To evaluate the maximum operating distance (Q6), the smartphones were positioned in front of the public display at a distance of 10 cm. The flashlight then emitted a specific test light signal 15 times in a row with a break of 1 second in between transmissions. The daemon process of the public display logged the recognized light signals to a file. At the end of a test, the experimenter counted the number of correctly (and incorrectly) recognized light signals, i.e., ideally 15 (or 0) signals. Afterwards, the distance between the camera and the smartphone was increased by 10 cm. The measurement was repeated until further increases did not change the observed results.

To evaluate the impact of different hardware configurations (Q7) on the number of correctly recognized light signals, each smartphone (D1–D5) was positioned 70 cm in front of the camera. This distance was chosen as it turned out to be half of the maximum operating distance (see below).

The test for the multi-user potential of the system (Q8) was conducted as follows: The smartphones D1 and D2 were positioned 70 cm in front of the camera and 30 cm apart from each other. Both smartphones then simultaneously emitted their unique light signal 15 times in a row. After counting the correctly and incorrectly recognized light signals per smartphone, the distance between the devices was reduced by 10 cm and the test was repeated until both were directly adjacent to each other.

Results

Lichtblick correctly recognized 11 to 15 light signals (73%–100%) if the smartphone was held at distances between 10 and 140 cm to the display. The number of incorrectly recognized light signals is between 0 and 2 (0%–13%). The system did not recognize any light signals (correct or incorrect) at distances greater than 140 cm. Figure 12.9 visualizes these results for the maximum operating distance (Q6). The counts do not add up to 15 if neither a correct nor an incorrect signal was detected. Table 12.12 summarizes the results for correctly and incorrectly recognized light signals emitted by each of the five smartphones (Q7): D1 and D2 provided the best recognition rates. The test results for the multi-user potential (Q8) are shown in Table 12.13: The detection rate of incorrectly interpreted signals is less or equal to 6%. This concludes the evaluation of visual interaction as to presenting the results obtained from the study. Subsection 14.2.3 is going to discuss the presented results.

Table 12.12.: Recognition of signals emitted by smartphone D1–D5.

	D1	D2	D3	D4	D5
Correct	80%	93%	33%	33%	53%
Incorrect	0%	0%	33%	20%	20%

Table 12.13.: Recognition of two signals (D1, D2) in parallel.

	30 cm	20 cm	10 cm	0 cm
Correct	86%, 86%	93%, 100%	86%, 100%	73%, 93%
Incorrect	0%, 6%	0%, 0%	0%, 0%	0%, 0%

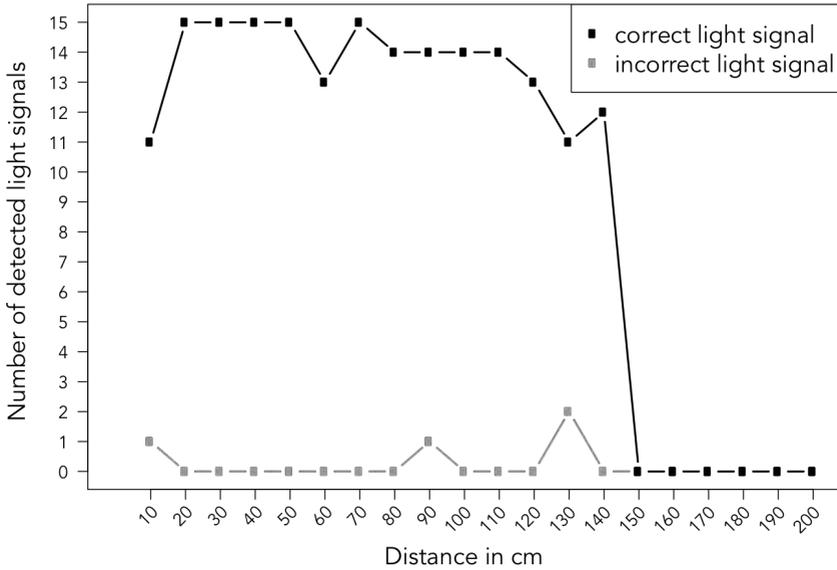


Figure 12.9.: Number of detected signals as a function of distance.

12.3. Process Integration

Though the IPED Toolkit and the Immersive Video Environment presented in Section 11.3 may not yet fully implement the approach proposed in Section 10.3, preliminary versions were used in a number of different contexts for initial evaluations. In addition, immersive environments have been successfully used in the past to evaluate various systems, see Subsection 10.3.2. This supports the idea of applying such an approach in the development process of public display systems. Finally, applying the proposed approach to an example scenario, as described on pp. 298, can also provide evidence for its usefulness in the context of developing public display systems. The following subsections evaluate both components of the process integration (C3) in more detail.

12.3.1. Immersive Public Display Evaluation and Design Toolkit

The evaluation of the IPED Toolkit is based on a long-term user study. This study was embedded in a seminar on “interactive public display systems” held in the summer term 2014. Nine students designed, developed, and evaluated a public display system. The students could freely choose an application scenario and should consider it with regard to the eight challenges introduced in Chapter 9. They were instructed to use the IPED Toolkit in all stages of the development process. At the end of the seminar, the students were asked to (anonymously) fill in a questionnaire to assess the suitability of the IPED Toolkit with regard to the eight challenges and perceived workload.

The questionnaire was divided into eight parts, one part for every challenge. Each part contained a NASA TLX and an *UMUX* question-

naire. The NASA TLX was expected to provide insights into the perceived mental, physical, and temporal demands, as well as the perceived performance, efforts, and frustration. The UMUX scores allow for identifying the challenges that the toolkit is more suitable for, and those challenges that the toolkit might be less suitable for. Finally, the students were asked to create screenshots of the virtual environment that they have created for their public display systems. At the end of the questionnaire, there was space for optional comments.

The students decided to design an indoor navigation and information system for the university. Users were able to search for lecture halls and seminar rooms in two ways: They could either enter a search phrase, e.g., the room number, or log into the interactive public display via their university account and select a course from their individual schedule. Once a lecture hall or seminar room was selected, the route from the users current position to the requested location was calculated and shown on an interactive map.

Table 12.14 presents the results of the NASA TLX questionnaires. The range of each NASA TLX scale reached from -10 (“very low”/“perfect”) to 10 (“very high”/“failure”), with zero as the neutral element. The individual results of all nine students were used to calculate the average score. The UMUX scores shown in Table 12.15 were calculated according to Finstad [78]. Some students also provided comments that support the proposed approach and its prototypical implementation:

The system may be used to quickly test various types of displays in varying situations. Testing occasions little costs.

—*Student 4 (English translation by the author of this thesis)*

Very supportive when designing a public display system.

—*Student 5 (English translation by the author of this thesis)*

Figure 12.10 shows three screenshots created by the students. As the questionnaires were filled in anonymously, it is not possible to relate the screenshots either to the comments presented above nor to the NASA TLX or UMUX scores. The background video footage was recorded by the students. It shows the entrance to the main lecture hall. There is a real public display is mounted on the wall next to the entrance, which can thus be seen in the recorded video footage. Figure 12.10a depicts the envisioned design of a public display mounted on the wall right next to the entrance. Figure 12.10b shows a public display, which is mounted on a special stand. This display is located underneath the public display, that already exists in the physical location. In contrast to the two designs presented above, Figure 12.10c includes static signage, i.e., a map and a banner, in addition to the public display system that was to be designed by the students. Just as the first example, all objects were placed next to the entrance door.

Table 12.14.: Results of the NASA TLX questionnaires (averaged scores) concerning the eight challenges introduced in Chapter 9 (M: mental demands, P: physical demands, PF: performance, E: efforts, F: frustration).

Challenge	M	P	T	PF	E	F
Situatedness	-2.11	-8.78	-0.56	-0.33	0.11	2.22
Form factors	-3.44	-9.33	0.78	-0.89	0.44	-0.33
Fixed e. f.	-3.11	-9.11	0.78	-2.33	-1.00	-1.89
Dynamic e. f.	-0.78	-8.89	0.56	-1.44	-1.78	-1.11
Mobile devices	-1.22	-7.67	0.44	-0.56	-1.11	0.33
Multi-display n.	-1.33	-8.11	-0.22	-0.89	-3.00	-0.33
Acceptance	-2.67	-8.67	-1.67	-4.33	-3.44	-2.11
Legal constraints	-1.44	-8.78	-1.11	-2.22	-2.11	-1.67



(a)



(b)



(c)

Figure 12.10.: Screenshots of the public display systems created by the students with the IPED Toolkit.

Table 12.15.: Results of the UMUX questionnaires concerning the eight challenges introduced in Chapter 9.

Challenge	UMUX score
Situatedness	55.56
Form factors	58.33
Fixed environmental factors	53.24
Dynamic environmental factors	51.39
Mobile devices	57.41
Multi-display networks	54.63
Acceptance	54.17
Legal constraints	54.63

This concludes the evaluation of the IPED Toolkit as to presenting the results obtained from the study. Subsection 14.3.1 is going to discuss the presented results.

12.3.2. Immersive Video Environment

Early versions of the IPED Toolkit and the IVE were demonstrated to different user groups on a number of occasions, e.g., a regional trade fair focused on smart cities. The demonstration included the gesture-based control mechanism described in Subsection 11.3.2. The mechanism enabled users to place various objects, including public displays, inside an urban scene that was depicted by captured video footage. While the interaction mechanism caused some issues, e.g., sensors not recognizing people and their gestures correctly, the general principle of augmenting video footage with objects to discuss, design, and experience public displays, for example, was easy to grasp for the participants. This is in line with observations from a series of lab-based

demonstrations with different audiences. The observations provide initial evidence that the approach proposed in this thesis is also accessible to non-experts. At another occasion, the IVE was demonstrated to a group of people working on a project about new digital signage that should be installed in the city of Münster. The members of the group had various backgrounds, from, e.g., urban planning to marketing, and can be considered experts in their respective fields. Overall, the reactions to the IVE were positive and encouraging, especially with regard to the perceived level of immersion. For instance, one person warned her colleagues about an approaching bus: “Watch out, the bus is coming!” (English translation by the author of this thesis). Another person uttered his excitement about the visual impression: “This allows to get a good impression about how it [the new digital signage system] would actually appear” (English translation by the author of this thesis). All members of the group walked around inside the IVE to experience the simulated environment from different perspectives.

Afterwards, a semi-structured interview, focusing on the six questions presented below, was conducted with all experts at the same time in a separate room. As the interview was administered in German, the questions and answers presented below were translated by the author of this thesis. The first question asked “how well the position of displays could be estimated and grasped” by using the Immersive Video Environment. “I think it is very realistic” was the reply of the first expert. The second expert agreed: “Really well. I think this is also interesting to the planners [of another department], since it’s always extremely difficult to work on models, as no one is able to imagine how it will eventually look like. I think it’s extremely important to use such a tool to visualize the design [of a project] and to include this visualization in political discussion, for example. I think this is way better than anything else we have today to visualize changes to the urban scenery.”

The second question asked whether it would be possible to “test and evaluate different form factors, e.g., the size, shape, or color, of such a system.” The first expert immediately replied: “Yes, absolutely. It was really impressive regarding the perceived perspective.” The second expert second this: “Absolutely.”

The third question tried to evaluate whether “the system would be suitable to represent networks, i.e., installations at different locations, of such systems.” In turn, the first expert wanted to know whether the system could show visual transitions between individual locations. As she was told that this is planned for a subsequent version, her answer was yes. The second expert elaborated this point a bit more: “If that was the case [referring to transitions], yes, because right now I think it’s hard to imagine that. Individual situations can be simulated well, but I don’t see how networks could be visualized.” The third expert added: “To do that, it is important to be able to move towards something and to see the perspectival changes.” In response to that, the first expert said: “Well, it would be possible to arrange the individual shots in such a way that the background of the first shot becomes the foreground in the second shot. This would help people to recognize that the drug store, which used to be in the background in the first shot, for example, is now closer to the user. This might help to understand the correlations between individual shots.”

The fourth question was about legal constraints: “Do you think that legal constraints could also be evaluated and assessed?” The second expert replied: “Yes, sure. I really think this is possible. For example, there are many issues related to traffic: ‘Would moving ads disturb drivers?’ or ‘where should this street sign be placed so that it can be seen by all drivers?’ It’s always about those questions.” The fourth expert noted that besides traffic, this would also apply to legal constraints from an artistic perspective.

“How well could this tool be integrated in existing workflows?” was the fifth question. “Well, it’s a visual inspection, which allows to see how something appears. One can draw further conclusions based on that.” responded the first expert. The second expert elaborated further: “I think it would be extremely helpful if we could say ‘we looked at this [scenario] and now go ahead and have a look at the situation [simulation].’ Especially with regard to the flow of traffic and pedestrians, such a system would be extremely useful. How would plazas look and feel like?” Building on this, the first expert added: “A possible extension would be to have something like a time line. This way, it would be possible to track and visualize the flow of pedestrians over a certain period of time, e.g., 24 hours. Such a time lapse could allow to identify all paths.” Once more, the interviewer asked whether the demonstrated system would actually be able to support or accelerate workflows and processes. The second expert stated: “Yes, absolutely, because [the system and its video footage] allows for a better perception than a plan. A plan may not consider the actual flow of people, as people don’t follow a plan. This allows to capture reality and thus to increase safety in traffic.”

Finally, the sixth question asked: “Was there something that disrupted you? Did you feel ill? Did you fail to recognize something? Is there something that might lessen the overall experience?” The first expert explained: “Well, I think the seams between the camera images appeared odd. They should be [more precisely] aligned to the edges of the projection screens. Besides that, I found it [the simulation] very realistic. It’s almost like being there. Probably, the sound also plays an important role. You should also be able to actually smell the exhaust gases [laughs].” The fourth expert asked: “Would it also be possible to have projections [overlays, e.g., simulated public displays] actually appear behind pedestrians, for example?” The fifth expert consented: “If there was something disturbing, then that. It appeared as if things

were simply pasted on top of everything.” The first expert agreed: “In the right moment, if there is nothing in front [of the overlay], it worked better. The question is, whether it would be acceptable to only use still imagery [instead of videos] and include some subtle movement. Of course, it wouldn’t be the same thing.”

At the end of the interview the first expert said: “I think it was really impressive that the displays could be moved and scaled while the video was playing in the background.” This was second by the sixth expert: “I also think it was really impressive, especially when this dove shot past me!”

13

Summary

This part focused on the design of privacy-preserving personalized public display systems. Chapter 9 introduced eight challenges that appear to have an impact on the design process of such systems. These challenges may serve as a first approach towards a model for public displays, which Davies et al. [62] identified to be missing. In turn, the approaches and prototypes presented in the subsequent chapters tried to address these design challenges.

Chapter 10 presented the theoretical approaches to the scientific contributions C1, C2, and C3 as introduced in Section 4.2. At first, Section 10.1 proposed the STRIDED* privacy threat model for public displays (C1), that is based on the renowned STRIDE model by Microsoft and the OWASP application security risks. The threat model helped to identify and prioritize privacy issues that personalized public display systems may be subject to: The results suggest that designers should focus on threats induced by spoofing and information disclosure first.

The next Section 10.2 set out to extend the list of countermeasures (C2) derived from an extensive literature review, as presented in Section 7.5. However, to address all identified threats would have exceeded the scope of this thesis significantly. Thus, this thesis focused

on three novel countermeasures that allow for bi-directional interaction between users and public displays: visual multiplexing, visual highlighting, and visual interaction. All three countermeasures use optical means of communication based on mobile devices, as the literature review indicates that using such a second device is the countermeasure used most frequently.

Finally, Section 10.3 presented an approach to consolidate the previously presented findings into a holistic process (C3). Researchers and designers of public display systems may build on this process and integrated it into existing structures. The approach consists of the IPED Toolkit and an Immersive Video Environment. Researchers and designers may use the toolkit to create simulations of public display systems with respect to the eight design challenges. The Immersive Video Environment, in turn, lets users experience the simulations created with the IPED Toolkit.

After the theoretical grounding to all approaches was laid out, Chapter 11 presented their prototypical implementations. The purpose of these prototypes was twofold: They should provide a solid basis for the subsequent evaluation as well as tangible scientific results (proof of concept). Thus, the privacy threat model was implemented as a public web application; the countermeasures visual multiplexing and visual highlighting are incorporated in a publicly available smartphone application that is accompanied by a corresponding web page; and finally, the IPED Toolkit became a software suite, which is available on a public source code hosting platform.

Based on these prototypes, Chapter 12 evaluated the approaches and their corresponding prototypes in user studies, surveys, and expert interviews. The next part of this thesis discusses the results obtained from the evaluation and draws final conclusions.

IV

Reflections

14

Discussion

This thesis provides three major scientific contributions as introduced in Section 4.2: (C1) a privacy threat model that researchers and designers of public displays can use to identify privacy issues; (C2) a set of countermeasures that address—at least some of—the issues identified by the model; and (C3) a methodology and tools that support the design, prototyping, and evaluation of privacy-preserving personalized public display systems. The preceding part presented the proposed approaches, that address the scientific contributions listed above. Once the theoretic foundation was laid out in Chapter 10, each approach was realized as a prototype, see Chapter 11. This prototype was then used to evaluate the individual approach, for example, via user studies in Chapter 12. This chapter reflects on these evaluations and discusses the results obtained. First, Section 14.1 considers the feedback on the privacy threat model gathered from the experts and the students. Afterwards, Section 14.2 examines the study results for the three countermeasures visual multiplexing, visual highlighting, and visual interaction. Section 14.3 then reviews the evaluations of the IPED Toolkit and the Immersive Video Environment—both parts of the process integration. Finally, Section 14.4 takes one step back and discusses all approaches and contributions as a whole.

14.1. Privacy Threat Model

The evaluation of the privacy threat model was based on a semi-structured expert interview, and a questionnaire filled in by students. The students also created threat models with the prototype presented in Section 11.1. Generally speaking, the answers provided by the experts as well as the students are promising; no expert pointed towards an apparent lack of threats, for example. “This is obviously a quick response from someone who has not really thought about it for a while, but it does look comprehensive.” Even though this comment by the fourth expert seems to be ad-lib, his expertise reinforces the significance and positive tenor of his statement. “The threat model is comprehensive enough now, but new threats may appear in the future.” This comment by the fifth student emphasizes an important fact: The threat model should not be considered as a static concept, that—once defined—will never change; it rather needs constant maintenance.

The comment of the second expert, however, may indicate an usability issue with the current prototype: The expert suggests that it should be possible to add individual threats to the threat model. In fact, users already may add unique items, e.g., threats, via the user interface, see pp. 264. Similarly, the first expert remarked, that it would be advisable to allow for combinations of threats or threat agents. This comment might also stem from a shortcoming in the user interface, as both, the theoretical model as well as the prototypical implementation, do not prevent users from adding arbitrary numbers of child nodes.

The same expert also commented that the meaning of the links between threat agents and threats remains unclear. This issue might be alleviated by adding labels to the links, similar to the labels shown in Figure 10.3. To better account for benign or malicious threat agents, for examples, the model could use variations of these labels, for ex-

ample, "...accidentally applying a ..." or "...willfully applying a ...". Furthermore, the first expert also speculates that there only is a limited set of threat models for all public displays. Based on this assumption, it would be an intriguing idea to compile a pool of threat models that designers may draw from. In the long run, this pool might help to design privacy-preserving public display systems. The second expert seconds this idea as he expects the privacy threat model to help designers consider privacy threats in advance, i.e., prior to deployment. In the same vein, the fourth expert remarks that all privacy threats might bear a certain resemblance. The envisioned pool of threat models would thus not necessarily grow unlimitedly.

Another important aspect is picked up by the second expert: "the spatial arrangement of the setting and the interaction." The relevance of this particular aspect is explained in Section 9.1. However, representing the situatedness in a (theoretical) model may be a challenging task, cf. the APEX framework presented in Subsection 7.6.1. Instead, this aspect is addressed by the IPED Toolkit, which allows to place public display systems within recorded video footage in a realistic manner. This way, such systems can be designed in an intuitive way.

With regard to supporting the design process of public displays, the experts attested the suitability of the approach to pursue this objective. The first and second expert, however, recommended to better guide users while using the prototype. One approach could be to lead users through a series of steps in a user interface. This concept is often referred to as a *wizard*. This wizard collects key data about a specific public display system and then automatically creates the rudiment of a corresponding threat model. The necessary information, i.e., the implied interrelations, could be drawn from the results of the literature survey as presented in Subsection 7.5.3. The same approach would also address some concerns uttered by the fourth expert.

The third and twelfth student emphasized the suitability of the system to support the design process in a “structured” and “systematic” way. These remarks indicate that the structure of the underlying concept is clear, reasonable, and comprehensible. The third expert emphasized the helpfulness of the tooltips guiding users through the design process. Certainly, students may tend to be less skeptical or sincere about tools they are supposed to use in a seminar, since they are concerned that genuine comments could be to their detriment. However, the students’ feedback still indicates a general—positive—tenor.

As described in Section 12.1, there are at least two approaches to creating privacy threat models for public displays with regard to a specific application scenario, i.e., breadth-first and depth-first. The former covers as many threat agents, threats, or weaknesses at once, while focusing on one level only: It does not consider possible threats that the chosen countermeasure may be subject to. The latter approach, in contrast, addresses one particular threat agent, agent, or weakness, in depth and considers at least one more level.

Both approaches may be valid and reasonable. Which one to choose probably depends on the individual application scenario at hand. For example, the breadth-first approach might be suitable when starting to design a public display system from scratch; breadth-first allows to quickly get an overview of all possible privacy threats. The depth-first approach, however, may be particularly suitable for subsequent development steps, for example, when improving an existing system for a specific application scenario. Either way, the threat models designed by the students turned out to be quite small, i.e., narrow as well as shallow. One reason for this could be the students’ limited knowledge and experience in the domain of privacy and security.

Finally, the results of the NASA TLX show that most perceived demands are low, i.e., below the neutral element. The only exception

are the perceived mental demands, see Table 12.1. This might indicate that designing a threat model may be perceived a complex task. This could also correlate with the students' knowledge and experience. To gain more insights into this, a subsequent study could compare the baseline, i.e., designing a threat model without the prototype, with the results obtained when using the prototype.

14.2. Countermeasures

This thesis proposed three countermeasures to address a number of privacy threats on public displays. Each countermeasure was implemented in a prototype, as presented in Section 11.2. This way, the proposed approaches can actually be applied by designers and researchers; at the same time, this also allows for a well-founded evaluation of the portrayed approaches. Section 12.2 reports on the results of that evaluation. The following subsections discuss the results with regard to each countermeasure: Subsection 14.2.1 focuses on visual multiplexing, Subsection 14.2.2 addresses visual highlighting, and Subsection 14.2.3 concentrates on visual interaction, eventually.

14.2.1. Visual Multiplexing

In general, users seem to be able to make use of visual multiplexing on public displays, since the overall positive study results were further backed up by the participants' feedback. For example, the sixth participant said: "Black and white images and texts were generally quite readable;" the fourteenth participant commented: "Finding the right answer was quite easy, except for the maps." Some participants even declared they would immediately install the software on their

mobile phones if such a system was actually deployed. Other participants proposed possible use case scenarios, in which the system could provide a certain level of privacy or security on public displays. In their opinion, the CDM approach could be extended to actually use different codes to encode individual information for specific audiences. Some participants actually enjoyed the tests and even challenged themselves to predict the correct answers without using the mobile device.

In summary, the outcomes could thus be interpreted as evidence for users being able to use visual multiplexing well (Q1, see Table 12.2). However, two comments were raised about the usability in terms of holding up a mobile device in front of your head in public: “I usually don’t walk around the city with my mobile at hand” (eighteenth participant); “The iPad may be too heavy for this purpose” (twentieth participant). One participant stated that the prototypical operating range between the public display and the mobile device was too small.

Q2 focused on the suitability of each multiplexing method for specific content types. H6 expected symbols to work best with any multiplexing method and the results support this. This might be due to the characteristics inherent to symbols, since they are designed to be easily recognizable and unambiguous. Their simple structures, large unicolored areas, and high contrasts facilitate good readability in most situations, even if partially obscured. This is probably also why FDM works best with this content type.

Though the characteristics of images showing words may seem similar to those that show symbols—and both should thus produce the same FDM results—the study revealed that there may be differences. One explanation could be that words have significantly smaller unicolored areas than symbols in most cases, cf. Figure 12.2c and 12.2d,

for example. Overall, the results for sentences are very comparable to those for words, see Figure 12.3.

In contrast to symbols, H7 expected the results for maps to turn out less satisfying. The study results supported this hypothesis. An explanation for this observation could be that—in most cases—maps hold considerably more information in the same area than symbols. For example, compare the details in Figure 12.2a to the details in Figure 12.2c. As explained in Subsection 10.2.1, FDM and CDM manipulate the input images in the process of generating a multiplexed image. As a result, the demultiplexed information channels may lack some details. TDM, however, does not alter the input images and thus retains their full information, which corresponds to the results of the study. The content type considered last were photographs. In this case, photos worked well with all analyzed multiplexing methods: 90.48% of all photos were recognized correctly.

While in theory, FDM should work as well as the other multiplexing methods, it turned out that in practice this was not the case. The results of the NASA TLX scores do not provide any striking information either: There are only minor differences in the median values, see Figure 12.6. However, the results of the test with a TV test card shed some light on the study results. Since the RGB values for green contain too many parts of blue, the demultiplexer erroneously assumes that parts of the green information channel belong to the blue channel. As a result, the information channels cannot be demultiplexed precisely and the individual images may interfere with each other.

The CDM results are considerably better than those for FDM for most content types, see Figure 12.4. Yet, the poor results for maps, cf. Figure 12.3, correspond to the outcomes of the NASA TLX scores: Participants reported high mental demands and efforts. They also tended to believe their answer was wrong and felt thus more frustrated than

they felt using any other multiplexing method. Especially with regard to the first prototypical implementation as used in the study, cf. pp. 267, CDM may have an impact on the quality of the demultiplexed image depending on the structure of the image: Horizontal patterns can be reconstructed well, whereas vertical patterns appear to be affected notably. This is probably why CDM performs poorly on maps, as they may contain dense horizontal as well as vertical patterns.

TDM returned good results overall, just a single participant identified a wrong demultiplexed map. All other content types were identified perfectly. This result thus provides strong support for H8. Figure 14.1 visualizes the superior image quality TDM produced. Motivated by the positive study results, the capabilities of the method were examined even further. Since TDM can be implemented to run in near realtime, a subsequent experiment used multimedia content, i.e., TV streams, as input material. Four different TV shows were captured and multiplexed with TDM, see Figure 11.7. Using the approach explained above, the demultiplexed videos ran in slow-motion due to the additionally injected video frames of the other information channels. Thus, when multiplexing four videos, plus the additional synchronization frame, their speed is reduced by a factor of five. In a second iteration, the objective was to preserve the original playback speed of each video. The multiplexer was programmed to only use every fifth frame of the input videos. The demultiplexed videos then actually ran at the desired speed, but motions appeared somewhat choppy since only every fifth frame was used. To overcome these issues, higher frame rates are required for displays as well as cameras embedded in mobile devices. At 100 fps, for example, the information channels, including the synchronization frame, could be transmitted very quickly. Each information channel could then be demultiplexed with 25 fps, which would ensure smooth playback.

The mental demand for all tested multiplexing methods appears to be reasonable, as do temporal demand, performance, effort, and frustration, see Figure 12.5 and Figure 12.6. There is neither a significant variance of means for each visual multiplexing method, nor for every content type. The system enables users to precisely judge their own performance, since the self-assessed NASA TLX performance corresponds to correctly respectively falsely given answers.

Recognizing demultiplexed maps imposed the highest demands on the participants. Especially when looking at maps multiplexed via CDM, participants reported feeling frustrated and having to invest quite some effort to complete the corresponding tasks. The NASA TLX based self-assessment of their performance also tends towards failure. In contrast to this, symbols could be easily recognized with low effort by all participants using any multiplexing method.

Quite unexpectedly, it turns out that CDM is slightly better in terms of NASA TLX scores than TDM. Comparing the performance of both methods, the opposite seems more likely: Even though TDM achieves 8.57 percentage points more than CDM, cf. Figure 12.4, participants ostensibly tend to prefer CDM. A closer look at the comments provides some insight: The flickering TDM video seems to be perceived as distracting and annoying. Coincidentally, the NASA TLX scores for TDM about spent efforts and the user's frustration are slightly higher than the ones for CDM, as depicted in Figure 12.6. However, participants felt most confident about their answer when using TDM, which can also be seen in Figure 12.6. In summary, the NASA TLX analysis showed that only very few participants consider the multiplexing methods as being frustrating and arduous; the majority felt successful and not frustrated while using the visual multiplexing prototype.



Figure 14.1.: Study participant using the TDM prototype. The public display currently shows a different video frame than the one selected by the user.

Limitations

Like any other study, the experiment carried out was subject to some limitations. The study was conducted indoors in a controlled lab environment, which avoided some issues facing public displays, e.g., reflections or glare. Nevertheless, many public displays are installed in indoor locations, and poor indoor results would have predicted even poorer outdoor results. Furthermore, only one specific monitor was used as a public display and one specific device as a mobile demultiplexer in order to limit the number of variables. The results might have varied if a range of screens and devices had been used. Mobile device with smaller screens may also impact the users' ability to recognize images. The results for each FDM information channel may also differ due to varying camera characteristics. The impact of hardware choices and combinations should thus be analyzed more thoroughly in subsequent studies.

The selection of content types was limited to those inspired by Schaeffler [195] and could be extended as well. The relatively small number of participants may also have an impact on the results of the study, particularly when looking at the ages and the participants' expertise with mobile phones.

Counter-balanced situations, e.g., based on latin square, would have allowed to test for sequence and training effects, as well as fatigue. Mixed linear models might be more suited than ANOVAs with regard to possible correlations between a participant's answers.

The study omitted scenarios in which the screen area is subdivided to suit the number of information channels, since there already is a large amount of research on this issue. This scenario would also contrast the presented motivation of reducing information overload, and has some inherent scalability problems.

Furthermore, the study did not investigate scenarios, in which information is delivered on a mobile device only, i.e., without public displays. Such scenarios would eschew the benefits of public displays, e.g., its situatedness, and suffer from scalability issues in very crowded settings, e.g., with respect to bandwidths or connectivity in general.

If frame rates of the camera could be significantly raised, more information channels could be transferred using the TDM approach. Also, the same number of information channels could be transferred less obtrusively. For example, instead of having time slots of equal length, there could be shorter and longer time slots. Depending on this length, there could be more salient as well as rather unobtrusive information channels. A TDM equipped public display operating at a speed of, e.g., 100 fps would reduce the length of each time slot to 0.01 seconds. The main information channel would be assigned 98 frames, so that it would be seen for 0.98 seconds. The remaining 0.02 seconds could be used for an additional information channel and the synchronization frame. Probably, users would hardly notice the flickering of the display and they could thus watch the main information channel with their bare eyes without interruption. Even if the flickering remains noticeable, it may render the display more interesting and thus alleviate display blindness [242]. Yet, the actual impact of such subliminal stimuli should be investigated, as they may still cause brain activity [39].

Currently, the FDM demultiplexer assumes a constant number of information channels and the TDM demultiplexer requires the number of information channels to be set manually. However, the TDM configuration could be automated in subsequent work, for example, by inferring the number of information channels based on the length and interval of the synchronization frame. The CDM approach already scales automatically, as the number of information channels is

conveyed in the QR tag. In terms of scalability, it should be noted once more that the amount of concurrent users is not limited, i.e., an infinite number of people can use the system at the same time.

After completing the study, the CDM prototype could be improved by using an optimized JPEG algorithm rather than using the naive implementation presented in Subsection 10.2.1, pp. 213. As this increases the visual quality of the demultiplexed images significantly, it would be interesting to repeat the user study with the new prototype. The outcomes may vary, especially with regard to maps, which performed poorly in the first study.

14.2.2. Visual Highlighting

The evaluation of the visual highlighting approach presented in Subsection 12.2.2 comprises two parts: (i) a comparison of the proposed approach with existing approaches, based a set of criteria introduced in Subsection 10.2.2; and (ii) an user study, that assessed the raw performance of visual highlighting with regard to efficiency, effectiveness, and robustness. The remainder of this subsection discusses the results of both parts.

Comparison of Existing Approaches

The analysis of visual highlighting methods on public displays reveals some shortcomings in four areas: (i) setup efforts, (ii) number of concurrent highlights, (iii) number of concurrent users, and (iv) imposed time constraints. While some of the presented work addresses these issues to a certain extend, none provides a solution to all of them.

In terms of setup efforts, CrossFlow and CrossBoard [43], for example, require a special software on the user's personal mobile device, that has been paired and synchronized to the public display in advance. Screen Codes [58] and SnapAndGrab [136] on the other hand do not require this previous synchronization, but are less functional for highlighting in return, i.e., they cannot point out the relevant information to the user. The approach proposed in this thesis and its prototypical implementation called Multipleye can be used on any public display with any type of hardware and software. The application for the personal mobile device can be used without prior network configuration or synchronization, i.e., no WiFi, 3G, or Bluetooth setup are required.

When looking at the number of concurrent highlights, CrossFlow and the Rotating Compass [189] can only handle one highlight at a time. Also the Interactive Ambient Public Displays [231] can only handle as many concurrent highlights as the number of concurrent users, which is assumed to be two or three. As for Multipleye, the number of concurrent highlights is limited by the size of the code tag and the resolution of the camera that is used to scan the code tag. Preliminary tests have shown that two to four concurrent highlights are feasible with current technologies, e.g., standing 2 m in front of a 24" display showing a code tag of 512 x 512 pixels.

Considering the number of concurrent users, Interactive Ambient Public Displays do offer highly personalized information, but can only handle two or three users simultaneously. Screen Codes and SnapAndGrab are more capable approaches in this respect as they can handle arbitrary numbers of users in parallel. Yet, Screen Codes can only carry one piece of information at a time and SnapAndGrab requires the user to actively select the desired information. Multipleye performs the actual visual highlighting on the user's personal mobile device, and does not require the public display to provide processing

power, bandwidth or any other kind of capacity. Thus, the number of concurrent Multipleye users is not limited by technical constraints.

Regarding time constraints, the reviewed approaches differ strongly. CrossFlow and the Rotating Compass require users to sync to the system first. Thus, it is not possible to use the system spontaneously. The duration of this synchronization process depends on the length of each highlighting cycle, which is in turn predefined by the total number of highlights and the time users need to sense and process the crossmodal cue. In contrast, Interactive Ambient Public Displays can be used instantaneously but can only serve a limited number of users at the same time. Multipleye can be used in a flexible and unrestricted manner. Firstly, users can aim their mobile devices at the public display to scan the visual tag at any time. Secondly, they can decide on how long the visual highlighting remains visible, as it is rendered on the personal mobile device and not on the public display itself. This also avoids possible interferences between different visual highlights or the actual information shown in the public display. Table 14.1 shows the benefits and drawbacks of the approach proposed in this thesis.

When visualizing the comparison between the systems using the spider diagram shown in Figure 12.7, a number of observations can be made. Firstly, the systems cover different points on the axes of several dimensions, e.g., granularity and concurrent highlights, whereas they are very similar in terms of concurrent users. Some dimensions, e.g., readability and interference, are less fully explored and indicate where further research might be needed.

User Study

The study did not analyze the effect of the participants' age on any of the measured variables, since the population was too small to gen-

erate reliable results. The reported learning effects in the first three sections are plausible, since the participants had to complete the particular tasks for the first time. This effect could have been avoided by showing the participants a number of training slides first or by permuting the order of appearance. Yet, the reported results are still valid, since the learning effect can easily be suppressed while not violating the integrity of the data.

Efficiency. According to the calculated efficiency values, visual highlighting improves the participants' performance significantly, especially when the content of the public display is visually distorted: 1.91 ($\bar{H}D$) in contrast to 3.80 (HD) is almost twice as efficient. Yet, according to the NASA TLX records for physical demands, the downside of the presented approach is that the participants have to hold up a personal mobile device. Whether the gain of efficiency justifies the additional physical efforts depends on the context the system is used in.

Effectiveness. Visual highlighting more than doubles the effectiveness of the participants: $1.68 (HD) / 0.70 (\bar{H}D) = 2.40$, cf. Table 12.8. Apparently, the stress that the participants experienced during the last four sections shifted their focus from physical demands towards efforts and frustration: Even though the participants had to hold up the same personal mobile device as before, the NASA TLX records for physical demands only show minor differences while the values for efforts almost tripled, cf. Table 12.10.

Workload. The analysis of NASA TLX scores does not reveal any inconsistencies. The participants felt lower mental demands, general

efforts, and frustration when using visual highlighting, while the confidence in their performance increased. However, the tablet caused a noticeable increase of the physical demands. It has yet to be decided whether the advantages outweigh the shortcomings. The use of mobile phones would likely reduce physical demands.

Robustness. In contrast to the assumptions, the applied visual distortions did not have a strong effect on the recorded results. There are statistical significant differences, but they are of minor relevance only, e.g., less than 1 second time difference, cf. Table 12.8. Overall, the system performed robustly and correctly in the given scenario.

Limitations

Though the study was designed to provide valid results, there may be some limitations to it. The study was conducted indoors, in a controlled lab environment, which may not be authentic for public displays. Still, the outcomes may be regarded as useful, as poor indoor results would have predicted even worse outdoor results. The study also used one specific display and one specific personal mobile device. In addition, the study focused on the raw performance of the system and the immediate impact on the user. Thus, it suppressed the time it takes users to take out and activate their personal mobile device. The participants' and the experimenter's reaction time may also have influenced the measured times. However, it can be assumed that this reaction time is relatively small compared to the overall measurement and can thus be neglected. Based on pre-tests, the number of wrong answers was neglected, due to their low occurrence. Finally, the relatively small sample of participants may also have had an impact on the results of the study.

Table 14.1.: Synopsis of the benefits and drawbacks of the visual highlighting prototype.

Benefits	Drawbacks
No conventional data connection, required; thus, no transmission charges and instant-on functionality.	The amount of transferable data is limited to the QR code specifications.
Depending on the situation, the best suited type of visual highlight can be chosen, e.g., rectangles, circles, and arrows.	–
Visual highlights work independently of each other and do not interfere with the conventional information on the public display since the highlighting is done on the user's personal mobile device.	The code tag takes up a certain amount of screen real estate.
Strong correlation between public display contents and visual highlights, as the code tags are directly embedded in the corresponding screens.	–
–	Users may be unaware of the advanced display capabilities, as QR tags are common on posters and bill boards.
The code tag does not change quickly over time, thus causing a calm visual impression.	This calmness restricts the number of concurrent highlights due to the limited amount of transferable data.

14.2.3. Visual Interaction

A lab-based study, presented in Subsection 12.2.3, evaluated the characteristics of the approach and the prototype. The study focused on the three questions Q6–Q8 as listed in Table 12.11. The remainder of this subsection addresses each question in more depth.

Maximum Operating Distance (Q6)

The evaluation of the baseline performance of Lichtblick provides evidence that the concept is a technical feasible approach to short-range optical interaction between public displays and smartphones. Users may use their smartphone to remotely control interactive public displays up to a distance of 140 cm. This range may be sufficient for many use cases, e.g., interacting with street maps. Yet, there may also be application scenarios which require users to stand farther away from the public display, e.g., when interacting with a very large media facade. However, it is very likely that this limitation is specific to the implementation of the prototype and the hardware used. It could thus be overcome with an optimized algorithm or more powerful hardware, e.g., cameras with higher resolutions.

Hardware Impacts (Q7)

The analysis shows that the overall recognition of the system varies considerably depending on the used smartphone. As mentioned in the Subsection 11.2.3, the fragmentation of the Android device market poses significant challenges when trying to precisely control the flashlight of the smartphone. One way to compensate for hardware and software inconsistencies could be to decrease the transfer speed

of Lichtblick adaptively until a communication could be established. Another way could be to increase the capture rate of the camera, which would decrease the transmission times of longer data streams in turn. The results of the IEEE Visible Light Communication Interest Group¹ show that higher speeds are achievable in principle. Either way, minimizing the number of incorrectly recognized light signals, i.e., false positives, would be highly desirable. The results of Q6 indicate that even the naive prototypical implementation can already handle transmission errors quite well, see Figure 12.9: only 13% of all tested light signals were interpreted erroneously, i.e., the public display performed an unwanted action. Additional logic, e.g., cyclic redundancy checks, could be added to avoid such misinterpretations.

Multi-User Potential (Q8)

Public displays are usually exposed to a broad audience. It is thus often necessary to handle multiple users in parallel. The test of the multi-user potential of Lichtblick provides some initial indication that the system scales well in setups with parallel users. The proposed approach could be regarded as an add-on to any third party software, that drives public displays. This third party software, in turn, would need to be able to handle simultaneous multi-user input reasonably well. It would thus be interesting to repeat the study based on real application scenarios and hardware setups. Moreover, the multi-user potential may be correlated to the actual application: Fast-paced games might have different timing requirements than map-based collaborations, for example.

Besides discussing the results of the three questions that guided the evaluation, further conclusions can be drawn. The remainder of this

¹<http://www.ieee802.org/15/pub/IGvlc.html>, accessed: March 6, 2015.

subsection discusses the outcomes of the study with regard to privacy and real-world deployments.

Privacy

The ability of Lichtblick to construct arbitrary user interfaces that are tailored to a specific application scenario provides designers of public displays great flexibility in terms of interaction. An interesting feature of Lichtblick is its ability to assign individual—and possibly secured—user interfaces to different users without having to identify the users. For example, two players of a game could receive control of different entities in the game that require different controls.

However, there may be privacy issues with regard to optical interaction between public displays and smartphones. The emitted light signals are visible to everyone in the vicinity. In principle, the communication between the smartphone and the public display could thus be overheard by a third party. Yet, the design of Lichtblick could be easily extended with existing encryption techniques, e.g., PGP, to counter such a threat: The public key of the display could be included in the QR code and the app could transfer its public key within the very first light signal. An even simpler approach could be to randomly change the payloads for each action of an action set. This would make it more difficult for attackers to reconstruct the triggered actions.

Besides the obvious blinking of the flashlight, the interaction between users and the public display is quite apparent, as users have to stand in close proximity to the displays while holding up their smartphone. This may hinder people to interact due to social inhibitions. Finally, the system—in particular its hardware—could be amended to use non-visible light, such as infrared (IR), to hide the interaction from onlookers. This would also reduce the risk of disturbing other people with

the blinking flashlight. Moreover, infrared light would address the privacy issue of capturing people in front of the public display: Both components of the system would be sensitive to invisible light only, so that faces, cars, or other object could not be traced. However, the special hardware requirements would conflict with the initial lightweight design principle.

Real-World Deployments

Reflective surfaces, such as window panes, may interfere with the optical interaction. A simple approach to filter these “ghost signals” would be to detect all light signals with very similar timings and to ignore all but the brightest one. Naturally, the lab-based evaluation did not provide any insights on such challenges that come with real-world deployments. Nevertheless, the results can be used as first insights into the general feasibility and as a baseline for further studies.

The latency between the emission of and the actual reaction to light signals may have an impact on the available action sets. The prototype requires 2.4 seconds to transmit one command, which might be a reasonable latency for selecting items via buttons, see Figure 10.12e (top), and adjusting values via sliders, see Figure 10.12e (bottom). However, controlling games via the d-pad, see Figure 10.12e (middle), may require more direct means of interaction.

Likewise, it would be interesting to investigate some usability aspects, such as fatigue: Holding up the smartphone for a prolonged period of time may become inconvenient. In turn, some application scenarios could be less suitable than others, e.g., playing a game, which requires constant interaction, compared to selecting items, which rather requires sporadic interaction.

It would thus be interesting to investigate whether other media, such as audio beyond the perceivable spectrum [162] would be more suitable. Higher transmission rates, e.g., 56 kbit/s might be achieved by using a technique similar to old modems. Additionally, users would be able to hold their smartphones in arbitrary positions, as no component, e.g., the flashlight, would have to face the public display. This could likely increase ergonomic aspects of the system.

Other application scenarios that may benefit from Lichtblick could be, for example: (i) Restaurants that let customers place orders via interactive public displays. Customers would only have to touch their own smartphone rather than the surface of the display before touching their food. (ii) Passengers in metro stations, who use Lichtblick to instantaneously interact with large entertainment displays, typically mounted out of reach behind the tracks. (iii) Hospitals could operate indoor navigation systems based on Lichtblick. This way, the system does not have to be touched and can thus poses no risk of transmitting germs via touch. As a novel approach to optical interaction between public displays and smartphones, Lichtblick thus offers interesting and unique features that provide a number of benefits in different settings.

Limitations

The study was carefully designed to provide valid results. Yet, there may be some limitations to it. For example, the study was conducted indoors, in a controlled lab environment, which may not be authentic for public displays. Still, the outcomes may be regarded as useful, as poor indoor results would have predicted even worse outdoor results. The study also used one specific monitor as the public display. Varying the used hardware components may lead to different results.

Furthermore, the prototype which was used in the study is based on a very naive image processing algorithm. An optimized version might yield better results in terms of recognition rates. In a similar vein, the camera which is used to receive the light signals of the smartphone could be replaced with a model that allows to capture images at higher frame rates. This could either be done to increase the transfer speed of the system, or to improve the recognition rate without influencing the speed. In both cases, using a higher resolution than 640 x 480 pixels likely provides better results.

Optimizing the binary codes, i.e., the payload, may also hold potential for improvement. The codes used in the study were quite long, i.e., 8 bit (see Table 11.1), which is a common size for primitive data types, for example, a *byte* in Java. The optical communication between the smartphone and the public display, however, may possibly work with even less bits, for example, a *nibble*, which consists of four bits. Optimizing the code by reducing the size each light signal may improve the recognition rate as well as the transmission speed. Finally, the relatively small sample of participants may also have had an impact on the results of the study.

14.3. Process Integration

Section 10.3 presented the scientific contribution of a process integration (C3), that incorporates the privacy threat model (C1), the countermeasures (C2), and the identified design challenges. This contribution consists of the IPED Toolkit and the Immersive Video Environment. The corresponding prototypes were presented in Section 11.3 and evaluated in Section 12.3. The following subsections discuss the results presented above.

14.3.1. Immersive Public Display Evaluation and Design Toolkit

Alt et al. [12] identified a number of research questions that appear to be of relevance in public display research. The approach presented in this thesis can help—at least to some degree—to address some of these questions related to user performance (e.g., task completion times or error rates), user experience (e.g., analyzing different interaction techniques), user acceptance (e.g., using a virtual prototype to support focus group discussions with a more realistic feel of the system), privacy (e.g., estimating threats such as shoulder surfing while entering data), and social impact (e.g., how to foster social interaction between strangers using the system for a specific task). The approach is less well suited to answer questions relating to audience behavior (requires real audiences in real settings) and display effectiveness (often assessed by observing people’s behavior in real settings).

The current IPED Toolkit implementation does not support 3D augmentation of video footage. It only allows for placing 2D representations of displays within the simulation. These 2D representations can be panned, scaled, rotated, and skewed to create the illusion of perspective and depth. A further limitation is the lack of sophisticated transitions between locations. Instead of “teleporting” users between recorded locations, smooth transitions, e.g., similar to Street View², might help users to create a mental map of the simulated area.

However, these drawbacks and limitations are outweighed by the benefits of the proposed approach. Compared to a design, prototyping, and evaluation process based on field studies, the approach would facilitate reproducibility while providing a high degree of visual realism. It thus offers a way to optimize the trade-off between internal

²<http://maps.google.com>, accessed: March 30, 2014.

and ecological validity. In the context of the local transportation scenario, for example, it would be interesting to analyze people's performance while using different versions of the system, such as different UI implementations, in a stressful situation, e.g., shortly before a bus arrives. While it would be difficult to repeatedly expose participants to this situation in a real environment, the proposed approach can easily facilitate this while immersing people in a realistic audiovisual simulation of the intended deployment site. At the same time, the effort required to carry out such a study is greatly reduced compared to a conventional field study: There is no need to transport people or equipment to study sites, for example.

While some aspects, such as appropriation or the impact of unforeseen factors, can only be fully assessed in field studies, the approach presented in this thesis can thus complement such studies in the ways described above, particularly at the early stages of the development when a functional version of a public display system is not available yet. Compared to lab-based studies, the approach increases the visual realism while providing means to record a variety of factors at the same level of detail.

The proposed approach also facilitates the integration and interaction with mobile devices by connecting the devices to a public display system via the state-transition graph (fifth challenge, Section 9.5). Developers can thus develop, test, and amend their software more rapidly as these steps can be carried out in the lab rather than at a (remote) deployment site. Multi-display setups can be simulated cost-efficiently, since the array of required devices is purely virtual (sixth challenge, Section 9.6). Furthermore, it is possible to change many characteristics, e.g., the form factor (second challenge, Section 9.2) or the position and rotation quickly and easily. This way, the (physical) effort as well as costs for development and user studies can be reduced.

Since the approach can simulate the appearance of public display systems realistically, even technically less savvy people, who are not involved in the actual development, can experience and use the system at early development stages prior to its actual installation. Designers can pinpoint possible design issues early on, e.g., by varying certain characteristics such as the display form factor (second challenge, Section 9.2) or by modifying certain environmental factors (third and fourth challenge, Sections 9.3 and 9.4). This may positively influence the user acceptance once the system is deployed. Similarly, designers might be able to assess legal constraints (eighth challenge, Section 9.8), e.g., legibility or distraction caused by displays.

Based on these considerations and the gathered experiences with using the approach, it can be inferred that it rather complements than replaces existing approaches to designing and evaluating public display systems. As previous work has pointed out [12], field studies are necessary to fully assess audience behavior, appropriation, or social impact, for example. Lab-based studies are very well suited to rigorously test hypotheses while exerting full control over a large number of variables. The proposed approach offers a middle-ground that combines aspects of both field as well as lab studies and may offer some key benefits particularly during the early development of public display systems. Initial experiences also suggest that this way of prototyping public display systems is accessible to designers and laypeople. In addition, the approach would lend itself well for a combination with model-driven approaches such as the ones introduced by Harrison and Massink [92] or Silva et al. [207]. Assessing the qualities and possibilities of such a combination, however, requires further studies.

Whether prototyping and deploying public displays systems in the virtual world is actually faster than a quick and dirty deployment at actual installation sites does not only depend on the target locations,

but also on the authoring tools available to designers. The current version of the IPED Toolkit provides only basic support and usability, so that the assessment of the efficiency of the proposed approach may be limited. Nevertheless, once deployment areas have been recorded and the corresponding graph has been created, it can be easily re-used to design and develop further public displays. For example, testing an alternative system to the public display for which the simulated environment was originally created, would simply require connecting the new public display system to the existing simulation.

The NASA TLX scores seem to speak out in favor of the IPED Toolkit. As explained in Subsection 12.3.1, zero represents the neutral element in the middle of each scale. With regard to the results, most scores tend towards a positive rating of the prototype: The perceived demands were low and the participants also felt successful in assessing the specific challenges. With regard to the perceived level of frustration, however, the situatedness of public display systems appears to be an exception. The NASA TLX score computes to 2.22, which is significantly above the neutral threshold. One reason might be that students felt frustrated about how overlays could be handled in the evaluated prototype. According to the comments gathered at the end of the questionnaire, the user interface seemed to be cumbersome, which might have led to poor results. However, the UMUX score for the same challenge contrasts the NASA TLX score: According to Table 12.15, situatedness is ranked third. This result may relativize the negative NASA TLX score to some extent.

The perceived mental demands for the second challenge, i.e., form factors, seem to be the lowest. This goes along with the highest UMUX score of 58.33. Apparently, the evaluated prototype actually helps users to design, prototype, and evaluate public display systems in a simulated environment with respect to size, shape, or color. In con-

trast to this, dynamic environmental factors appear to require the most perceived mental demands. This observation is further backed up by the lowest UMUX score of 51.39. This outcome may have been expectable, as certain factors, such as social inhibition, may not be simulated in a lab-based environment, as explained at the beginning of this subsection. With regard to the NASA TLX results for perceived performance, the challenges of acceptance and fixed environmental factors may be assessed the best. While this outcome speaks out in favor of the concept of the approach, the corresponding UMUX scores indicate that the prototypical implementation still bears some potential for improvement.

Limitations

As with any user study, this experiment may have been subject to some limitations. First of all, students may tend to be less skeptical or sincere about tools they are supposed to use in a seminar, since they are concerned that genuine comments could be to their detriment. However, the students' feedback still indicates a general tendency. Furthermore, the small sample size of nine students may have had an impact on the overall outcome. It would thus be interesting to repeat the study with more participants or with participants that have another social background, e.g., experts in urban planning. Finally, the results may also be tied to the chosen application scenario, i.e., an indoor navigation and information system for the university campus. There probably are individual characteristics to every application scenario. Thus, future work could systematically analyze different application scenarios and compare the obtained results.

14.3.2. Immersive Video Environment

Based on experiences gathered so far, specific drawbacks and advantages of the Immersive Video Environment can be identified. One key limitation is the lack of support for locomotion. This is an inherent problem of virtual environments and has been a research subject for a long time [55, 197]. While omnidirectional treadmills are a still very expensive way to address this, using photos or videos to construct simulations further limits user movement, as only the recorded views can be experienced without distorting images. Consequently, the approach proposed here is better suited to investigate scenarios where locomotion is not essential.

In terms of privacy analysis, the proposed approach can be used to carry out controlled studies, e.g., simulating shoulder surfing in a specific situation. It is, however, not well suited to assess audience behavior, display effectiveness, or social impact as those aspects heavily depend on various characteristics of the actual installation site.

Furthermore, the prototypical implementation suffers from a number of limitations. For example, the movement of the mirror image avatar is not restricted, so that users can place it in physically impossible positions, e.g., floating above ground; this could break the immersion. Finally, both the avatar and virtual objects are simply overlaid over the video footage: Moving objects such as cars that intersect with these simply disappear behind them regardless of where they are supposed to be in the 3D space defined by the video. This is another aspect that can negatively affect immersion.

Most of these limitations can be addressed by improving the current implementation. A more sophisticated and robust gesture recognition system would allow for rotation gestures and enable multi-user interaction. Video scenes recorded by a moving camera constitute a

more difficult problem that is not easily overcome. One option would be to have a dynamic layer model that changes over time as the camera moves; this might, however, complicate interaction, e.g., objects disappearing during interaction, and also induce nausea.

A more sophisticated layer model could also specify permissible locations of the avatar on each layer to prevent avatars from being moved to physically impossible locations. Realizing physically correct occlusions involving the avatar would require a deeper analysis of the video and a more sophisticated spatial model.

Finally, several limitations relate to the used gestures. While the gestures were learned quickly by the participants and positively received in the user study, further studies are required to identify the most immersive or intuitive set of gestures. So far, only a subset of all the gestures, i.e., movement control, was assessed. In addition, it was not tested whether the use of devices, e.g., mobile phones, to carry out certain actions, such as injecting virtual objects, would be more immersive or intuitive—neither on their own nor in combination with gestures. Further studies on these aspects are desirable as well.

Generally speaking, mirror image avatars could also be used with photographs or 3D virtual environments. Using video footage of the user was expected to let the real-time motion of the avatar blend better with the genuine movement occurring in the recorded video footage; it would thus create a strong sense of presence and immersion. Using 3D virtual environments would allow for correct occlusions but constructing realistic virtual worlds requires a lot of effort. Compared to a desktop scenario, in which users would place objects and experience augmented video scenes, it can be argued that the gesture-based approach combined with a large screen provides a more realistic and immersive experience. Initial informal feedback from people seeing the system in action as well as observations from the initial user study

on the movement gestures seem to confirm this. Yet, it would be advisable to carry out a series of user studies to investigate these aspects in more detail. Further studies are also needed to identify the most suitable gesture sets for different tasks.

According to the experts' answers gathered in the semi-structured interview, see Subsection 12.3.2, the situatedness of public displays can be simulated quite well in the Immersive Video Environment. The second expert points out that the perceived experience may serve as the basis for discussion with other project partners in a way that cannot be achieved when using conventional plans or sketches. The experts also attested the appropriateness of the Immersive Video Environment to design, prototype, and evaluate the form factor as well as legal aspects of public display systems.

As already mentioned, smooth transitions between individual locations appear to be a key aspect, that should be investigated further in the subsequent development. The first three experts agreed that multi-display networks could only be addressed adequately if users were able to understand the transitions and thus build a mental model of the simulated display network. There are multiple approaches to this, for example, by using a technique similar to Street View (see above) or recording pre-defined transitions. The latter one would yield better visual results, while the user's degree of freedom might be limited, as it is not likely that all possible transitions could be taken into consideration in advance.

The experts agree that the proposed approach and prototype could be integrated well into existing workflows. Once more, the experts emphasize that the proposed approach could be used as a basis for discussion with other people, such as citizens. The prototype was envisioned to become a powerful visualization tool, here to analyze traffic flows, that could be used to facilitate citizen engagement. However, the ex-

perts also said that the visual representation should be optimized in order to provide a realistic impression of the simulated environment. This could be achieved by using specialized camera equipment rather than the DIY approach presented in Subsection 11.3.2, for example.

14.4. Summary

Each of the previous sections focused on one scientific contribution (C1–C3) and discussed it in detail. This section takes a step back and looks at all contributions from a higher perspective. Overall, privacy gained interest in society as well as in science in the last years, see p. 9 and p. 119. At the same time, public displays proliferated in urban environments and became a ubiquitous experience. However, some of the advantages offered by this technology lie fallow, since people tend to actively ignore public displays—the phenomenon is referred to as display blindness. One cause for this blindness towards public displays is the lack of relevant content. Personalized content is regarded as relevant and may thus be suited to address the root of this phenomenon and the negative effects for all stakeholders. Yet, personal content on public displays calls for certain means of privacy in turn. It is thus timely and reasonable to address the design of privacy-preserving personalized public display systems in this thesis.

Even though privacy constitutes a major aspect of personalized public display systems, some issues lay beyond the scope of this thesis. For example, further research could focus on how to actually implement and perform personalization, or how to acquire and interpret data for personal profiles. A further interesting strand of research could analyze people's reactions to privacy breaches: What would be the consequences if private or sensitive information became public unintentionally? What would be people's reactions, with regard to the

user whose information has been disclosed as well as bystanders? In the long term, new social conventions—or transmission principles, according to Nissenbaum—could emerge.

From a similar perspective, it would be interesting to analyze whether and how society would adapt to personalized information on public display systems. The following scenario is an example: Nowadays, non-personalized public displays are used to inform travelers about arriving, departing, and delayed trains at train stations. All travelers waiting at the same platform receive the same information about a delayed train simultaneously. Reactions of individual persons, such as sighing or cursing, are thus comprehensible for others. If the same information would have been communicated through a privacy-preserving personalized public display, however, the individual reaction might not be as reasonable and the public could be confused. A similar effect could be observed when people started to use barely noticeable headsets to place calls with their mobile phones. Bystanders were led to think that the person soliloquizes. Further elaboration on such intriguing aspects is clearly beyond the scope of this thesis. Thus, the remainder of this section focuses on the scientific contributions C1–C3 with regard to privacy.

The privacy threat model (C1) is a novel approach to systematically assess and evaluate privacy issues for public display systems. It is based on established concepts, for example, STRIDE, OWASP, or the stakeholders proposed by Alt et al. [5, 7, 8], see Subsection 2.2.1. This renders the model robust as the underlying components are actually used in production environments: STRIDE has been used by Microsoft for many years. In comparison to alternatives such as EBIOS, SP 800-30, or OCTAVE (see Subsection 10.1.1), the proposed threat model is more compact, concrete, and comprehensible. Moreover, it is not based on commercial products and may thus be used for free. This way, it sup-

ports designers, researchers, as well as laypeople in the complex task of designing privacy-preserving personalized public display systems. For this purpose, the theoretical model was also implemented as a public prototype. Eventually, the privacy threat model also allows for a systematic comparison of different systems or approaches along a unified set of criteria. In this regard, the privacy threat model constitutes a major scientific contribution. A shortcoming of such a privacy threat model could be that it constraints people in the design process: People might be tempted to focus on the items covered by the model only and avoid “thinking out of the box,” considering, e.g., individual threats. Instead, the model should be regarded as a starting point that creates a common basis for further discussion and evaluation.

Privacy is an essential requirement for the personalization of public displays; personalization is the key to relevant content on public displays; and relevant content, in turn, is expected to pave the way for accepted and successful public displays. The list of existing countermeasures and the three novel approaches as proposed in this thesis thus constitute a significant scientific contribution (C2). Based on an extensive literature survey, the existing countermeasures were sighted, clustered, and finally condensed to a useable form. The corresponding heat map further helps to identify applicable countermeasures and to compare different approaches systematically. Besides providing a common ground for discussion, the list of countermeasures and the heat map thus support designers in creating actual privacy-preserving systems. Still, there are also possible issue with regard to C2. For example, the list of countermeasures could become incomplete and might require updating. Furthermore, the heat map could be inaccurate, i.e., it suggests a less well suited or even an inappropriate countermeasure. As with the privacy threat model, users—designers and researchers—might be tempted to refer to the list or heat map unreflectingly.

Nevertheless, the heat map also revealed an apparent paradox: Public display systems tend to avoid showing personalized content on the public display itself. On the contrary, this information is often shown on second devices, such as smartphones instead. This either indicates that public displays are inappropriate for personalized content per se, or that research mainly avoided this sensitive issue so far. The three novel countermeasures address this paradox and further broaden the spectrum of available means for privacy: The personalized content is actually shown on the public display, but a second device, i.e., a smartphone, provides access to it. The three countermeasures were designed as generic means, as they are not tailored to a specific application scenario, cf. Table 7.2. Along with their prototypical implementations, the novel countermeasures thus constitute another considerable scientific contribution of this thesis. However, the study results indicate that some are more suitable for certain application scenarios than others. This observation is likely caused by current technological constraints, such as frame rates, for example. It might thus be interesting to harness upcoming technological advances in future work.

Finally, Davies et al. [62] note that research on public display systems is still in its infancy. Thus, there is a lack of established methods, tools, and techniques that researchers may reach for. The three contributions presented in this thesis strive to compensate for this lack. Integrating the proposed methodology and tools into a holistic process (C3) may foster the advent of new application areas, cf. [62, pp. 92–93]. The Immersive Video Environment allows multiple stakeholders, including laypeople, to simultaneously design and experience public display systems in a high-fidelity simulation. This characteristic fosters interdisciplinary communication and the exchange of ideas. In comparison to deploying prototypes in situ, the approach proposed in this thesis allows for rapid prototyping at significantly lower costs.

The mirror image avatar further helps to reduce interaction barriers and entice people to interact. Though this sounds promising and beneficial, future work should evaluate the prerequisites that need to be met so that the general public—laypeople in particular—can use the methodology proposed in this thesis in a sensible way: A certain level of privacy-awareness might be necessary in order to understand the underlying concepts and to use the proposed methods and tools to somebody's advantage. Today's society, for example, seems to be privacy-agnostic at best: The "I've got nothing to hide" argument [213] appears in most public discussions about privacy at some point. It might thus be desirable to establish a profound societal understanding of privacy as early as possible. One approach would be, for example, to interweave this sensitive issue in school education.

There is also room for improvements, for example, with regard to transitions or locomotion. The first one could help people to create a mental map of the simulated environment while wandering around. The latter could increase the sense of immersion for application scenarios in which locomotion is of importance. Along the same lines, it should be pointed out that the Immersive Video Environment may not be capable of simulating all relevant environmental aspects: The audience behavior or the social surrounding in general can likely be assessed *in situ* only. Though C3 can be applied and used as it is, combining it with other frameworks, e.g., P-LAYERS, may provide additional input or guidance for designers and researchers. Future work could analyze which frameworks are compatible and supplemental. The individual components of C3 could benefit from a community-driven development, since the source codes are publicly available³. Thus, the software components as well as the underlying concepts can be extended continuously based on a collaborative process.

³<https://github.com/sitcomlab/IPED-Toolkit>, accessed: July 15, 2015

15

Conclusion

As explained in Section 4.2, this thesis provides the three main contributions C1–C3. Each contribution is summarized in the remainder of this section, together with further contributions that emerged from the work and results presented in this thesis. Table 15.1 presents an overview of all research questions and the corresponding scientific contributions. Table 15.2 summarizes the practical contributions of this thesis. Moreover, future work that may be inspired by the outcomes presented in this thesis or that might exceed the scope of this thesis is also discussed.

15.1. Contributions

This thesis evaluated the applicability of the STRIDE threat model to public displays in terms of privacy. The study results indicate that STRIDE can be used to model major privacy threats. However, the meaning of the letter D should be changed from denial of service to decontextualization. Thus, the modified threat model can be used for future analysis of interactive public displays with regard to privacy. The results also identified the relative importance of these privacy threats. There is an apparent discrepancy between the participants'

explicit and implicit prioritizations of those threat categories. Yet, the results imply that public display engineers should especially focus on privacy threats induced by either information disclosure or spoofing. The privacy threat model thus constitutes the scientific contribution *C1* to the first research question.

The results presented in the context of *C1* also define a design space for privacy demands on public displays. Engineers of public displays can use this design space to build privacy-aware public display systems that align with users' privacy perceptions and needs more closely. Results from a preliminary qualitative evaluation indicate that the threat model is comprehensive and supports the design and engineering process of privacy-aware interactive systems.

Based on an extensive literature survey, this thesis compiled a list of 25 existing countermeasures that can be used to address—at least some—of the privacy issues identified by *C1* above. This list of countermeasures, which constitutes *C2*, is also represented as a heat map, that allows designers and researchers of public display systems to quickly identify the most commonly used countermeasure for a particular privacy threat. Additionally, this thesis presented and evaluated three novel countermeasures, i.e., visual multiplexing, visual highlighting, and visual interaction, that add to the list of countermeasures (*C2*).

Visual multiplexing comprises three techniques: frequency-division multiplexing (FDM), code-division multiplexing (CDM), and finally time-division multiplexing (TDM). Using a prototypical implementation realizing all three methods, a user study was carried out to contrast the approaches when applied to five different types of content. The results indicate that participants were able to use the system successfully with little training and manageable workload. There seem to be differences between the three multiplexing methods in terms of success rate and suitability for different content types. Of the five

content types, symbols worked best and maps showed the worst performance. Overall, the results provide initial evidence that all three visual multiplexing methods enable multiple users to concurrently access personalized multimedia content in a privacy-preserving way on public displays.

Existing approaches to visual highlighting on public displays suffer from one or more of four issues: (i) setup efforts, (ii) number of concurrent highlights, (iii) number of concurrent users, and (iv) imposed time constraints. This thesis introduced an approach to address these issues. The approach uses composite visual tags and personal mobile devices to visually highlight personally relevant information. The thesis also proposed a set of comparison criteria, which can be used to compare visual highlighting on public displays and to further explore the design space. The results of a user study found that the approach to visual highlighting has genuine potential to help users locate information on public displays in a privacy-preserving way. The results showed a significant gain of speed, up to twice as fast as conventional systems. The analysis of the efficiency of the system indicates that user may benefit from visual highlighting in situations in which time is of importance. The results overall suggest that participants were able to use the system successfully with little training and manageable workload. The system appears to be robust and work correctly even in situations with visual distortions on the public display.

This thesis also presented a novel approach for privacy-preserving interaction between public displays and smartphones based on short-ranged optical communication. The overall feasibility of the approach was demonstrated by evaluating a prototype in a lab-based study. In comparison to existing approaches, the system offers a number of advantages, e.g., bidirectional communication that is independent from external radio-based network infrastructures, e.g., Bluetooth or WiFi.

The approach thus avoids additional costs, especially roaming fees, and setup overheads, e.g., pairing processes. Users may thus use the system in a direct and instantaneous manner. Assuming a secure way of processing the camera images, tracking users is also harder in comparison to network-based approaches since there are no unique identifiers, e.g., IP addresses, for the users' devices. Using the approach, public displays can offer interactivity without exposing easily breakable input devices, e.g., keyboards, that could be the victim of vandalism or potentially unhygienic. Due to the lightweight hardware and software requirements, the proposed approach can be easily and cost efficiently applied to existing public display installations. The system can be used in multi-user environments and supports a broad range of application scenarios due to dynamically defined user interfaces.

Moreover, this thesis proposed a novel approach to engineer public display systems based on realistic audiovisual simulations and a state-transition graph. The approach has been integrated in a holistic process, which constitutes *C3*. Researchers and designers of public displays can directly apply this process. In this context, the key contributions include a systematic analysis of approaches to engineer public display systems, a novel approach that integrates many of the benefits of previous approaches, an architecture for a toolkit implementing the approach, and an initial assessment of the approach based on an example scenario and first experiences from using the prototype. Key benefits of the proposed approach include high re-usability of simulated environments, reduced effort to construct deployment sites and scenarios, as well as support for a broad range of prototypes, e.g., of varying fidelity, and design and evaluation methods. This work can thus contribute towards simplifying and accelerating the development of privacy-preserving public display systems.

This thesis also proposed a novel approach to interact with video environments in an immersive and intuitive way. Using an avatar, users can move inside the footage in three dimensions and place virtual objects inside the video scene. Knowledge about the height of the user and the layer model enable the system to place the video avatar in three dimensions. Besides privacy-preserving public displays, the system can be used for various applications, for example, the prototyping and evaluation of ubiquitous and situated systems: An IVE could be used as a means to prototype augmented reality applications. The initial prototype—though limited—used web technologies to illustrate the feasibility of the approach. A study provided initial evidence for a high degree of immersion and the usability of the proposed approach.

Finally, this thesis identified six research opportunities, which are also based on the outcome of the literature survey. Future work in the domain of privacy-preserving public displays could thus focus on these threats and application scenarios by Perry and O'Hara [183]: tampering with regard to “social grooming,” “current and past working processes,” and “planning and information overview”(O1–O3); as well as decontextualization with regard to “demonstrating achievements,” “current and past working processes,” and “planning and information overview” (O4–O6).

Table 15.1.: Scientific contributions to the research questions that guided this thesis.

ID	Research question and scientific contributions
RQ1	<p data-bbox="288 400 972 424">What are main privacy threats on public displays?</p> <p data-bbox="288 443 1012 671">This thesis proposes a privacy threat model (C1) based on the STRIDED* model that addresses most privacy issues and also prioritizes the privacy threats: information disclosure and spoofing should be addressed first. Moreover, C1 also spans a design space for privacy-sensitive application scenarios: Reading personal messages, browsing photos, and using social networks requires the most privacy on public displays.</p>
RQ2	<p data-bbox="288 699 975 722">What are countermeasures to those privacy threats?</p> <p data-bbox="288 742 1012 1002">This thesis provides a classification of existing countermeasures in literature (C2). This classification is based on a list obtained from an extensive literature survey. In addition to the existing countermeasures, this thesis proposed three novel approaches to facilitate privacy-preserving, bi-directional communication between users and public displays: (i) visual multiplexing, (ii) visual highlighting, and (iii) visual interaction. With regard to visual highlighting, this thesis also contributes a set of comparison criteria.</p>
RQ3	<p data-bbox="288 1029 986 1053">How to support the design process of public displays?</p> <p data-bbox="288 1072 1012 1305">This thesis proposes a methodology to design, prototype, and evaluate privacy-preserving personalized public display systems (C3): a novel method to engineer such systems based on realistic audiovisual simulations and a state-transition graph; a systematic analysis of approaches to engineer public displays; and an architecture for a toolkit; common design challenges that may be used to guide the design process of public display systems.</p>

Table 15.2.: Practical contributions to the research questions that guided this thesis.

ID	Research question and practical contributions
RQ1	<p data-bbox="197 515 874 539">What are main privacy threats on public displays?</p> <p data-bbox="197 560 916 616">This thesis provides the privacy threat model (C1) as a publicly available prototype.</p>
RQ2	<p data-bbox="197 639 874 663">What are countermeasures to those privacy threats?</p> <p data-bbox="197 684 916 914">This thesis presents a heat map of countermeasures, which is based on the classification of countermeasures (C2). Designers and researchers of privacy-preserving personalized public display systems can use this heat map for identify the most commonly used countermeasures with regard to a specific privacy threat. Furthermore, the novel countermeasures of visual multiplexing and visual highlighting have been implemented in publicly available prototypes.</p>
RQ3	<p data-bbox="197 938 891 962">How to support the design process of public displays?</p> <p data-bbox="197 983 916 1182">This thesis realizes realistic audiovisual simulations of public display systems based on an Immersive Video Environment (IVE). To allow for natural interaction with this simulation, this thesis implemented and evaluated the concept of mirror image avatars. The architecture of the underlying toolkit has been implemented as the publicly available IPED Toolkit.</p>

15.2. Future Work

A phenomenon frequently observed in science is that providing an answer to a particular question leads to more questions in turn. Consequently, based on the scientific contributions provided by this thesis, new research potentials may open up. Future work may draw from these potentials to push the knowledge in the domain of privacy-preserving public display systems even further. First of all, subsequent work could address the six research opportunities as identified above. Similarly, it would be desirable to update the literature survey with recent publications, in order to keep track of the most recent advances. Furthermore, applying a different classification scheme, e.g., the “space of input device and display possibilities” as proposed by Dix and Sas [73], to the same body of work could yield further interesting results in terms of privacy on personalized public display systems.

With regard to the privacy threat model and its prototype, the results of the literature survey could be used to realize a wizard that guides users through the design process of a privacy-preserving public display system. It would also be interesting to repeat the user study with a baseline, i.e., the perceived mental demands when not using the model and prototype proposed in this thesis. Additionally, future work could look into countermeasures and focus on evaluating the proposed threat model in real world scenarios more thoroughly. The web-based tool, which provides CI in a tangible way, could also be extended, improved, and evaluated.

Based on the results on visual multiplexing, several promising areas for further research emerged. A logical next step would be to evaluate the approach in the real world rather than in the lab, and to compare it with solutions that rely solely on mobile phones, such as in navigation scenarios, for example. Another aspect worthy of further investi-

gation is to only multiplex parts of the screen and leave the remaining areas untouched. This would enable anyone to access information on the screen, while providing those with demultiplexing devices with personalized content. A further interesting line of research relates to non-destructive visual multiplexing methods. There, additional information is embedded without rendering the screen content unreadable to the human eye. The concept of steganography would be comparable to this approach. The collaborative newspaper by Lander et al. [124] is another example. Finally, there is room to optimize the multiplexing methods from a technical perspective.

Subsequent work on visual highlighting could analyze the overall performance of the system instead of the raw performance as evaluated in this thesis. It could be investigated how to reduce the size of the QR tag and how to increase the operating range. An updated version of the QR code technology, called iQR codes, may allow for multiple parallel visual highlights in a single iQR tag. This way, the presented approach could provide a number of individual personalized visual highlights. Finally, analyzing different types of visual highlights, e.g., magnifying glasses or fish eyes instead of colored rectangles, may open up interesting research areas.

Based on the encouraging results on visual interaction, several areas for further research emerged: As the prototypical implementation presented in this thesis is quite naive, it would be interesting to incorporate more robust approaches to optical data transfer [94, 147]. Based on this enhanced prototype, it would be intriguing to analyze real-world deployments in varying application scenarios. One particular aspect would be to explore the potential of adaptive (XUL-based) user interfaces and to evaluate varying action sets in different application scenarios. Future research could also focus on how to increase transfer speeds, e.g., based on concepts by Haas and colleagues [202].

In addition to expanding the existing version of the IPED Toolkit, there are a number of further directions for future research. One of them relates to simulating movement within the IVE, e.g., via a treadmill and footage of movement, and with extending the field of view of the IVE, i.e., by using head-mounted displays. Closely linked to this aspect is the investigation of different ways to visualize movement, e.g., via different cinematographic transitions or via footage of actual movement. A further interesting line of work concerns the way in which different people can interact with the system, e.g., groups of designers discussing alternatives or end-users providing input during the early stages of a participatory design. Finally, and possibly most importantly, there is a need to systematically compare different design and evaluation methods for public display systems via controlled user studies. In particular, it would be very valuable to clearly establish similarities and differences between field studies and studies carried out in the IVE, e.g., with respect to quantifying the impact of different contextual factors. The IPED Toolkit and the Immersive Video Environment could also be evaluated by more users with a broader range of expertise in varying application scenarios, for example at design workshops with multiple stakeholders. Similar user studies could also explore the characteristics of mirror image avatars as a means of interaction in simulated environments. Similar studies could compare the avatars to alternative input means, e.g., 3D controllers. Furthermore, it could be investigated how mobile devices could be integrated into the system, e.g., as a secondary controller to let users select content.

Eventually, it would be intriguing to integrate all the proposed approaches and countermeasures in a single holistic tool. This way, the design process of privacy-preserving public display systems could be streamlined in a well-integrated, universal process.

V

Appendix

Bibliography

- [1] IEEE Standard Glossary of Software Engineering Terminology. *IEEE Std 610.12-1990* (1990).
- [2] ACM SIGCHI Curricula for Human-computer Interaction. Tech. rep., ACM, 1992. ISBN: 0-89791-474-0.
- [3] Digital signage market—forecast to 2014–2020. Tech. rep., MarketsandMarkets, 2014. <http://bit.ly/12sur7w>, accessed: July 14, 2015.
- [4] Ahn, S., Lee, T.-S., Kim, I.-J., Kwon, Y.-M., and Kim, H.-G. Large Display Interaction Using Video Avatar and Hand Gesture Recognition. In *Image Analysis and Recognition*, A. Campilho and M. Kamel, Eds., vol. 3211 of *LNCS*. Springer, 2004.
- [5] Alt, F. *A design space for pervasive advertising on public displays*. PhD thesis, Universität Stuttgart, 2013.
- [6] Alt, F., Balz, M., Kristes, S., Shirazi, A. S., Mennenöh, J., Schmidt, A., Schröder, H., and Goedicke, M. Adaptive User Profiles in Pervasive Advertising Environments. In *Proc. Aml '09*, Springer (2009).
- [7] Alt, F., Kubitzka, T., Bial, D., Zaidan, F., Ortel, M., Zurmaar, B., Lewen, T., Shirazi, A. S., and Schmidt, A. Digifieds: insights into deploying digital public notice areas in the wild. In *Proc. MUM '11*, ACM (2011).

- [8] Alt, F., Memarovic, N., Elhart, I., Bial, D., Schmidt, A., Langheinrich, M., Harboe, G., Huang, E., and Scipioni, M. Designing Shared Public Display Networks – Implications from Today’s Paper-Based Notice Areas. In *Pervasive Computing*, K. Lyons, J. Hightower, and E. Huang, Eds., vol. 6696 of *LNCS*. Springer, 2011.
- [9] Alt, F., Memarovic, N., Greis, M., and Henze, N. UniDisplay – A research prototype to investigate expectations towards public display applications. In *PERCOM ’14 Workshop*, IEEE (2014).
- [10] Alt, F., Müller, J., and Schmidt, A. Advertising on Public Display Networks. *Computer* 45, 5 (2012).
- [11] Alt, F., Schneegass, S., Girgis, M., and Schmidt, A. Cognitive effects of interactive public display applications. In *Proc. PerDis ’13*, ACM (2013).
- [12] Alt, F., Schneegaß, S., Schmidt, A., Müller, J., and Memarovic, N. How to evaluate public displays. In *Proc. PerDis ’12*, ACM (2012).
- [13] American Psychological Association, Ed. *Publication Manual of the American Psychological Association*, 6 ed. American Psychological Association, 2009.
- [14] Andrés del Valle, A. C., and Opalach, A. The Persuasive Mirror: Computerized Persuasion for Healthy Living. In *Proc. HCI Int. ’05* (2005).
- [15] Ardito, C., Buono, P., Costabile, M. F., and Desolda, G. Interaction with Large Displays: A Survey. *ACM Comput. Surv.* 47, 3 (2015).
- [16] Ardito, C., Costabile, M. F., Lanzilotti, R., De Angeli, A., and Desolda, G. A Field Study of a Multi-touch Display at a Conference. In *Proc. AVI ’12*, ACM (2012).

-
- [17] Azad, A., Ruiz, J., Vogel, D., Hancock, M., and Lank, E. Territoriality and behaviour on and around large vertical publicly-shared displays. In *Proc. DIS '12*, ACM (2012).
- [18] Azuma, R. T. A survey of augmented reality. *Presence: Teleoperators and Virtual Environments* 6, 4 (1997).
- [19] Baldauf, M., Fröhlich, P., and Lasinger, K. A scalable framework for markerless camera-based smartphone interaction with large public displays. In *Proc. PerDis '12*, ACM (2012).
- [20] Baldauf, M., Lasinger, K., and Fröhlich, P. Private public screens: detached multi-user interaction with large displays through mobile augmented reality. In *Proc. MUM '12*, ACM (2012).
- [21] Baldauf, M., Salo, M., Suetterle, S., and Fröhlich, P. Display Pointing: A Qualitative Study on a Recent Screen Pairing Technique for Smartphones. In *Proc. BCS-HCI '13*, British Computer Society (2013).
- [22] Baldauf, M., Suetterle, S., Fröhlich, P., and Lehner, U. Interactive Opinion Polls on Public Displays: Studying Privacy Requirements in the Wild. In *Proc. MobileHCI '14*, ACM (2014).
- [23] Ballagas, R., Rohs, M., and Sheridan, J. G. Sweep and point and shoot: phonecam-based interactions for large public displays. In *Adj. Proc. CHI '05*, ACM (2005).
- [24] Ballendat, T., Marquardt, N., and Greenberg, S. Proxemic interaction: designing for a proximity and orientation-aware environment. In *Proc. ITS '10*, ACM (2010).
- [25] Benko, H., and Wilson, A. D. Multi-point interactions with immersive omnidirectional visualizations in a dome. In *Proc. ITS '10*, ACM (2010).

- [26] Berger, S., Kjeldsen, R., Narayanaswami, C., Pinhanez, C., Podlaseck, M., and Raghunath, M. Using Symbiotic Displays to View Sensitive Information in Public. In *Proc. PERCOM '05* (2005).
- [27] Beyer, G., Alt, F., Müller, J., Schmidt, A., Isakovic, K., Klose, S., Schiewe, M., and Hauelsen, I. Audience behavior around large interactive cylindrical screens. In *Proc. CHI '11*, ACM (2011).
- [28] Bier, E. A., Stone, M. C., Pier, K., Buxton, W., and DeRose, T. D. Toolglass and magic lenses: the see-through interface. In *Proc. SIGGRAPH '93*, ACM (1993).
- [29] Boehm, B. A Spiral Model of Software Development and Enhancement. *SIGSOFT Softw. Eng. Notes* 11, 4 (1986).
- [30] Böhmer, M., Gehring, S., Löchtefeld, M., Ostkamp, M., and Bauer, G. The Mighty Un-touchables – Creating Playful Engagement on Media Façades. In *Adj. Proc. MobileHCI '11* (2011).
- [31] Böhmer, M., Hecht, B., Schöning, J., Krüger, A., and Bauer, G. Falling Asleep with Angry Birds, Facebook and Kindle: A Large Scale Study on Mobile Application Usage. In *Proc. MobileHCI '11*, ACM (2011).
- [32] Böhmer, M., and Müller, J. Users' Opinions on Public Displays that Aim to Increase Social Cohesion. In *Proc. IE '10*, IEEE (2010).
- [33] Boring, S., Baur, D., Butz, A., Gustafson, S., and Baudisch, P. Touch projector: mobile interaction through video. In *Proc. CHI '10*, ACM (2010).
- [34] Boring, S., Gehring, S., Wiethoff, A., Blöckner, A. M., Schöning, J., and Butz, A. Multi-user interaction on media façades through live video on mobile devices. In *Proc. CHI '11*, ACM (2011).

- [35] Boring, S., Jurmu, M., and Butz, A. Scroll, tilt or move it: using mobile phones to continuously control pointers on large public displays. In *Proc. OzCHI '09*, ACM (2009).
- [36] Bouzit, M., Burdea, G., Popescu, G., and Boian, R. The Rutgers Master II-new design force-feedback glove. *IEEE/ASME Transactions on Mechatronics* 7, 2 (2002).
- [37] Boyle, M., and Greenberg, S. The language of privacy: Learning from video media space analysis and design. *TOCHI* 12, 2 (2005).
- [38] Brignull, H., and Rogers, Y. Enticing people to interact with large public displays in public spaces. In *Proc. INTERACT '03* (2003).
- [39] Brooks, S. J. a. o. Exposure to subliminal arousing stimuli induces robust activation in the amygdala, hippocampus, anterior cingulate, insular cortex and primary visual cortex: A systematic meta-analysis of fMRI studies. *NeuroImage* 59, 3 (2012).
- [40] Brudy, F., Ledo, D., Greenberg, S., and Butz, A. Is Anyone Looking? Mitigating Shoulder Surfing on Public Displays Through Awareness and Protection. In *Proc. PerDis '14*, ACM (2014).
- [41] Burke, M., Hornof, A., Nilsen, E., and Gorman, N. High-cost Banner Blindness: Ads Increase Perceived Workload, Hinder Visual Search, and Are Forgotten. *TOCHI* 12, 4 (2005).
- [42] Calderon, R., Blackstock, M., Lea, R., Fels, S., de Oliveira Bueno, A., and Anacleto, J. RED: A Framework for Prototyping Multi-display Applications Using Web Technologies. In *Proc. PerDis '14*, ACM (2014).

- [43] Cao, H., Olivier, P., and Jackson, D. Enhancing Privacy in Public Spaces Through Crossmodal Displays. *Social Science Computer Review* 26, 1 (2008).
- [44] Cardoso, J., and José, R. A Framework for Context-Aware Adaptation in Public Displays. In *OTM '09 Workshop*, R. Meersman, P. Herrero, and T. Dillon, Eds., vol. 5872 of *LNCS*. Springer, 2009.
- [45] Cardoso, J., and José, R. PuReWidgets: a programming toolkit for interactive public display applications. In *Proc. EICS '12*, ACM (2012).
- [46] Carsten, T. P., Röcker, C., Streitz, N., Stenzel, R., and Magerkurth, C. Hello.Wall – Beyond Ambient Displays. In *Adj. Proc. Ubicomp '03* (2003).
- [47] Caudell, T. P., and Mizell, D. W. Augmented reality: an application of heads-up display technology to manual manufacturing processes. In *Hawaii Int. Conf. on System Sciences* (1992).
- [48] Chang, K. S.-P., Danis, C. M., and Farrell, R. G. Lunch Line: Using Public Displays and Mobile Devices to Encourage Healthy Eating in an Organization. In *Adj. Proc. UbiComp '14*, ACM (2014).
- [49] Chen, X. A., Boring, S., Carpendale, S., Tang, A., and Greenberg, S. Spalendar: Visualizing a Group's Calendar Events over a Geographic Space on a Public Display. In *Proc. AVI '12*, ACM (2012).
- [50] Cheverst, K., Clarke, K., Fitton, D., Rouncefield, M., Crabtree, A., and Hemmings, T. SPAM on the Menu: The Practical Use of Remote Messaging in Community Care. In *Proc. CUU '03*, ACM (2003).
- [51] Cheverst, K., Dix, A., Fitton, D., Kray, C., Rouncefield, M., Sas, C., Salsis-Lagoudakis, G., and Sheridan, J. G. Exploring bluetooth

- based mobile phone interaction with the hermes photo display. In *Proc. MobileHCI '05* (2005).
- [52] Cheverst, K., Dix, A., Fitton, D., Kray, C., Rouncefield, M., Saslis-Lagoudakis, G., and Sheridan, J. Exploring Mobile Phone Interaction with Situated Displays. *PERMID, Pervasive '05 Workshop* (2005).
- [53] Chignell, M. H., Quan-Haase, A., and Gwizdka, J. The Privacy Attitudes Questionnaire (PAQ): Initial Development and Validation. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 47, 11 (2003).
- [54] Child, S. “Car” from The Noun Project. <http://thenoun-project.com/term/car/1597>, accessed: May 19, 2015.
- [55] Chim, J., Lau, R. W. H., Leong, H. V., and Si, A. CyberWalk: a web-based distributed virtual walkthrough environment. *IEEE Transactions on Multimedia* 5, 4 (2003).
- [56] Churchill, E., Girgensohn, A., Nelson, L., and Lee, A. Blending digital and physical spaces for ubiquitous community participation. *Commun. ACM* 47, 2 (2004).
- [57] Churchill, E. F., Nelson, L., and Hsieh, G. Café life in the digital age: augmenting information flow in a café-work-entertainment space. In *Adj. Proc. CHI '06*, ACM (2006).
- [58] Collomosse, J. P., and Kindberg, T. Screen codes: visual hyperlinks for displays. In *Proc. HotMobile '08*, ACM (2008).
- [59] Conny and Conan. Spiral model (Boehm, 1988). [http://commons.wikimedia.org/wiki/File:Spiral_model_\(Boehm,_1988\).svg](http://commons.wikimedia.org/wiki/File:Spiral_model_(Boehm,_1988).svg), accessed: February 11, 2015.

- [60] Cox, D., Kindratenko, V., and Pointer, D. IntelliBadge: Towards Providing Location-Aware Value-Added Services at Academic Conferences. In *UbiComp 2003: Ubiquitous Computing*, A. Dey, A. Schmidt, and J. McCarthy, Eds., vol. 2864 of LNCS. Springer, 2003.
- [61] Dang, C. T., and André, E. A Framework for the Development of Multi-display Environment Applications Supporting Interactive Real-time Portals. In *Proc. EICS '14*, ACM (2014).
- [62] Davies, N., Clinch, S., and Alt, F. Pervasive Displays: Understanding the Future of Digital Signage. *Synthesis Lectures on Mobile and Pervasive Computing* 8, 1 (2014).
- [63] Davies, N., Friday, A., Newman, P., Rutledge, S., and Storz, O. Using Bluetooth Device Names to Support Interaction in Smart Environments. In *Proc. MobiSys '09*, ACM (2009).
- [64] Davies, N., Langheinrich, M., Clinch, S., Elhart, I., Friday, A., Kubitza, T., and Surajbali, B. Personalisation and Privacy in Future Pervasive Display Networks. In *Proc. CHI '14*, ACM (2014).
- [65] Davies, N., Langheinrich, M., Jose, R., and Schmidt, A. Open Display Networks: A Communications Medium for the 21st Century. *Computer* 45, 5 (2012).
- [66] Davis, F. D. Perceived Usefulness, Perceived Ease of Use, and User Acceptance of Information Technology. *MIS Q.* 13, 3 (1989).
- [67] De Luca, A., and Frauendienst, B. A Privacy-respectful Input Method for Public Terminals. In *Proc. NordiCHI '08*, ACM (2008).
- [68] De Luca, A., Langheinrich, M., and Hussmann, H. Towards Understanding ATM Security: A Field Study of Real World ATM Use. In *Proc. SOUPS '10*, ACM (2010).

- [69] Dearman, D., and Truong, K. N. BlueTone: a framework for interacting with public displays using dual-tone multi-frequency through bluetooth. In *Proc. Ubicomp '09*, ACM (2009).
- [70] Delikostidis, I., Fechner, T., Fritze, H., AbdelMouty, A. M., and Kray, C. Evaluating Mobile Applications in Virtual Environments: A Survey. *Int. J. Mobile Human Computer Interaction (IJMHCI)* 5, 4 (2013).
- [71] Delikostidis, I., Fritze, H., Fechner, T., and Kray, C. Bridging the Gap Between Field- and Lab-Based User Studies for Location-Based Services. In *Progress in Location-Based Services 2014*, G. Gartner and H. Huang, Eds., Lecture Notes in Geoinformation and Cartography. Springer, 2015.
- [72] Dey, A. K. Understanding and Using Context. *Personal Ubiquitous Comput.* 5, 1 (2001).
- [73] Dix, A., and Sas, C. Mobile Personal Devices meet Situated Public Displays: Synergies and Opportunities. *Int. J. Ubiquitous Computing (IJUC)* 1, 1 (2010).
- [74] Elhart, I., Langheinrich, M., Memarovic, N., and Heikkinen, T. Scheduling Interactive and Concurrently Running Applications in Pervasive Display Networks. In *Proc. PerDis '14*, ACM (2014).
- [75] Faisal, T., and Cheverst, K. Exploring User Preferences for Indoor Navigation Support through a Combination of Mobile and Fixed Displays. In *Proc. MobileHCI '11*, ACM (2011).
- [76] Farnham, S. D., McCarthy, J. F., Patel, Y., Ahuja, S., Norman, D., Hazlewood, W. R., and Lind, J. Measuring the Impact of Third Place Attachment on the Adoption of a Place-based Community Technology. In *Proc. CHI '09*, ACM (2009).

- [77] Federal Trade Commission. Security Check: Reducing Risks to your Computer Systems. Tech. rep., Washington, 2003. <https://www.ftc.gov/system/files/documents/plain-language/bus58-security-check-reducing-risks-your-computer-systems.pdf>, accessed: May 4, 2015.
- [78] Finstad, K. The Usability Metric for User Experience. *Interacting with Computers* 22, 5 (2010).
- [79] Friday, A., Davies, N., and Efstratiou, C. Reflections on Long-Term Experiments with Public Displays. *Computer* 45, 5 (2012).
- [80] Gamma, E., Helm, R., Johnson, R. E., and Vlissides, J. *Design Patterns. Elements of Reusable Object-Oriented Software.*, 1st ed., reprint. ed. Prentice Hall, Reading, Mass, 1994.
- [81] Geel, M., Huguenin, D., and Norrie, M. C. PresiShare: opportunistic sharing and presentation of content using public displays and QR codes. In *Proc. PerDis '13*, ACM (2013).
- [82] Gehring, S., and Krüger, A. Facade map: continuous interaction with media facades using cartographic map projections. In *Proc. UbiComp '12*, ACM (2012).
- [83] Gehring, S., and Wiethoff, A. Interaction with Media Façades. *Informatik-Spektrum* (2014).
- [84] German Federal Office for Information Security. IT-Grundschutz Catalogues: Version 2005. Tech. rep., 2005. https://www.bsi.bund.de/EN/Topics/ITGrundschutz/itgrundschutz_node.html, accessed: July 14, 2015.
- [85] Gotardo, P. F. U., and Price, A. Integrated space: authoring in an immersive environment with 3d body tracking. In *SIGGRAPH '10 Posters*, ACM (2010).

- [86] Greenberg, S., Boring, S., Vermeulen, J., and Dostal, J. Dark Patterns in Proxemic Interactions: A Critical Perspective. In *Proc. DIS '14*, ACM (2014).
- [87] Greenberg, S., and Rounding, M. The Notification Collage: Posting Information to Public and Personal Displays. In *Proc. CHI '01*, ACM (2001).
- [88] Hall, E. T. *The Hidden Dimension*. Anchor Books, 1990.
- [89] Hamhoun, F., and Kray, C. Scalable navigation support for crowds: personalized guidance via augmented signage. In *Proc. COSIT '11*, Springer (2011).
- [90] Hamhoun, F., and Kray, C. Supporting pilgrims in navigating densely crowded religious sites. *Personal and Ubiquitous Computing* 16, 8 (2012).
- [91] Handte, M., Wagner, S., Apolinarski, W., and Marron, P. J. iScreen: a toolkit for interactive screens. In *Proc. PerDis '12*, ACM (2012).
- [92] Harrison, M. D., and Massink, M. Modelling Interactive Experience, Function and Performance in Ubiquitous Systems. *Electronic Notes in Theoretical Computer Science* 261 (2010).
- [93] Hernan, S., Lambert, S., Ostwald, T., and Shostack, A. Uncover security design flaws using the STRIDE approach. *MSDN Magazine* (2006). <http://msdn.microsoft.com/en-us/magazine/cc163519.aspx>.
- [94] Hesselmann, T., Henze, N., and Boll, S. FlashLight: optical communication between mobile phones and interactive tabletops. In *Proc. ITS '10*, ACM (2010).

- [95] Horak, R. *Webster's New World Telecom Dictionary*, 1 ed. Webster's New World, 2007.
- [96] Hosio, S. *Leveraging Social Networking Services on Multipurpose Public Displays*. PhD thesis, University of Oulu, 2014.
- [97] Hosio, S., Kukka, H., and Riekkki, J. Social Surroundings: Bridging the Virtual and Physical Divide. *MultiMedia, IEEE 17*, 2 (2010).
- [98] Huang, E., Koster, A., and Borchers, J. Overcoming Assumptions and Uncovering Practices: When Does the Public Really Look at Public Displays? In *Pervasive Computing*, J. Indulska, D. Patterson, T. Rodden, and M. Ott, Eds., vol. 5013 of *LNCS*. Springer, 2008.
- [99] Huang, E. M., and Mynatt, E. D. Semi-public displays for small, co-located groups. In *Proc. CHI '03*, ACM (2003).
- [100] Huang, E. M., Mynatt, E. D., and Trimble, J. P. When design just isn't enough: the unanticipated challenges of the real world for large collaborative displays. *Personal Ubiquitous Comput. 11*, 7 (2007).
- [101] Huang, E. M., Russell, D. M., and Sue, A. E. IM here: public instant messaging on large, shared displays for workgroup interactions. In *Proc. CHI '04*, ACM (2004).
- [102] ITS. Multiplexing; Frequency-division multiplexing; Code-division multiple access; Time-division multiplexing. www.its.bldrdoc.gov/fs-1037/dir-023/_3439.htm, www.its.bldrdoc.gov/fs-1037/dir-016/_2344.htm, www.its.bldrdoc.gov/fs-1037/dir-007/_1033.htm, www.its.bldrdoc.gov/fs-1037/dir-037/_5453.htm, accessed: Augsut 22, 2012.

- [103] Izadi, S., Brignull, H., Rodden, T., Rogers, Y., and Underwood, M. Dynamo: a public interactive surface supporting the cooperative sharing and exchange of media. In *Proc. UIST '03*, ACM (2003).
- [104] Izadi, S., Kim, D., Hilliges, O., Molyneaux, D., Newcombe, R., Kohli, P., Shotton, J., Hodges, S., Freeman, D., Davison, A., and Fitzgibbon, A. KinectFusion: real-time 3d reconstruction and interaction using a moving depth camera. In *Proc. UIST '11*, ACM (2011).
- [105] Jancke, G., Venolia, G. D., Grudin, J., Cadiz, J. J., and Gupta, A. Linking Public Spaces: Technical and Social Issues. In *Proc. CHI '01*, ACM (2001).
- [106] José, R., Cardoso, J., Alt, F., Clinch, S., and Davies, N. Mobile applications for open display networks: common design considerations. In *Proc. PerDis '13*, ACM (2013).
- [107] José, R., Otero, N., Izadi, S., and Harper, R. Instant Places: Using Bluetooth for Situated Interaction in Public Displays. *IEEE Pervasive Computing* 7, 4 (2008).
- [108] José, R., Pinto, H., Silva, B., Melro, A., and Rodrigues, H. Beyond Interaction: Tools and Practices for Situated Publication in Display Networks. In *Proc. PerDis '12*, ACM (2012).
- [109] Jurmu, M. *Towards engaging multipurpose public displays: design space and case studies*. PhD thesis, University of Oulu, 2014.
- [110] Kamijo, K., Kamijo, N., and Gang, Z. Invisible barcode with optimized error correction. In *Proc. ICIP '08* (2008).
- [111] Kaur, N., and Kaur, P. Mitigation of SQL Injection Attacks Using Threat Modeling. *SIGSOFT Softw. Eng. Notes* 39, 6 (2014).

- [112] Kaviani, N., Finke, M., Fels, S., Lea, R., and Wang, H. What goes where?: designing interactive large public display applications for mobile device interaction. In *Proc. ICIMCS '09*, ACM (2009).
- [113] Kim, D., Hilliges, O., Izadi, S., Butler, A. D., Chen, J., Oikonomidis, I., and Olivier, P. Digits: freehand 3d interactions anywhere using a wrist-worn gloveless sensor. In *Proc. UIST '12*, ACM (2012).
- [114] Kim, S., Cao, X., Zhang, H., and Tan, D. Enabling concurrent dual views on common LCD screens. In *Proc. CHI '12*, ACM (2012).
- [115] Kjeldskov, J., and Paay, J. A longitudinal review of Mobile HCI research methods. In *Proc. MobileHCI '12*, ACM (2012).
- [116] Kohnfelder, L., and Garg, P. The threats to our products. *Interface (internal Microsoft journal)* (1999). <http://blogs.msdn.com/b/sdl/archive/2009/08/27/the-threats-to-our-products.aspx>, accessed: May 15, 2015.
- [117] Kray, C., Cheverst, K., Fitton, D., Sas, C., Patterson, J., Rouncefield, M., and Stahl, C. Sharing control of dispersed situated displays between nand residential users. In *Proc. MobileHCI '06*, ACM (2006).
- [118] Kray, C., and Delikostidis, I. Evaluating location-based services (a position paper). In *GeoHCI, CHI '13 Workshop* (2013).
- [119] Kray, C., Kortuem, G., and Krüger, A. Adaptive Navigation Support with Public Displays. In *Proc. IUI '05*, ACM (2005).
- [120] Kray, C., Nesbitt, D., Dawson, J., and Rohs, M. User-defined gestures for connecting mobile phones, public displays, and table-tops. In *Proc. MobileHCI '10*, ACM (2010).

- [121] Kurdyukova, E., Bee, K., and André, E. Friend or foe? relationship-based adaptation on public displays. In *Proc. AmI'11*, Springer (2011).
- [122] Kurdyukova, E., Hammer, S., and André, E. Personalization of Content on Public Displays Driven by the Recognition of Group Context. In *Ambient Intelligence*, F. Paternò, B. Ruyter, P. Markopoulos, C. Santoro, E. Loenen, and K. Luyten, Eds., vol. 7683 of *LNCS*. Springer, 2012.
- [123] Kurdyukova, E., Obaid, M., and André, E. Direct, bodily or mobile interaction?: comparing interaction techniques for personalized public displays. In *Proc. MUM '12*, ACM (2012).
- [124] Lander, C., Speicher, M., Paradowski, D., Coenen, N., Biewer, S., and Krüger, A. Collaborative Newspaper: Exploring an adaptive Scrolling Algorithm in a Multi-user Reading Scenario. In *Proc. PerDis '15*, ACM (2015).
- [125] Langheinrich, M. Privacy in Ubiquitous Computing. In *Ubiquitous Computing*, J. Krumm, Ed. CRC Press, 2009.
- [126] Langheinrich, M., Schmidt, A., Davies, N., and José, R. A practical framework for ethics: the PD-net approach to supporting ethics compliance in public display studies. In *Proc. PerDis '13*, ACM (2013).
- [127] Lee, C., Bonebrake, S., Hollerer, T., and Bowman, D. A. The Role of Latency in the Validity of AR Simulation. In *Proc. VR '10*, IEEE (2010).
- [128] Lee, K., Clinch, S., Winstanley, C., and Davies, N. I Love My Display: Combatting Display Blindness with Emotional Attachment. In *Proc. PerDis '14*, ACM (2014).

- [129] Li, M., Arning, K., Sack, O., Park, J., Kim, M.-H., Ziefle, M., and Kobbelt, L. Evaluation of a Mobile Projector-Based Indoor Navigation Interface. *Interacting with Computers* (2013).
- [130] Lindén, T., Heikkinen, T., Ojala, T., Kukka, H., and Jurmu, M. Web-based framework for spatiotemporal screen real estate management of interactive public displays. In *Proc. WWW '10*, ACM (2010).
- [131] Lucero, A., Holopainen, J., and Jokela, T. MobiComics: collaborative use of mobile phones and large displays for public expression. In *Proc. MobileHCI '12*, ACM (2012).
- [132] Lyle, P., Lueg, C., and Nugent, T. Multi-cursor Multi-user Mobile Interaction with a Large Shared Display. In *Proc. OzCHI '12*, ACM (2012).
- [133] Marquardt, N., Diaz-Marino, R., Boring, S., and Greenberg, S. The proximity toolkit: prototyping proxemic interactions in ubiquitous computing ecologies. In *Proc. UIST '11*, ACM (2011).
- [134] Marshall, P., Morris, R., Rogers, Y., Kreitmayer, S., and Davies, M. Rethinking 'Multi-user': An In-the-wild Study of How Groups Approach a Walk-up-and-use Tabletop Interface. In *Proc. CHI '11*, ACM (2011).
- [135] Matusik, W., Forlines, C., and Pfister, H. Multiview user interfaces with an automultiscopic display. In *Proc. AVI '08*, ACM (2008).
- [136] Maunder, A., Marsden, G., and Harper, R. Creating and sharing multi-media packages using large situated public displays and mobile phones. In *Proc. MobileHCI '07*, ACM (2007).

- [137] Maunder, A. J., Marsden, G., and Harper, R. SnapAndGrab: accessing and sharing contextual multi-media content using bluetooth enabled camera phones and large situated displays. In *Adj. Proc. CHI '08*, ACM (2008).
- [138] McCarthy, J. F., Congleton, B., and Harper, F. M. The Context, Content & Community Collage: Sharing Personal Digital Media in the Physical Workplace. In *Proc. CSCW '08*, ACM (2008).
- [139] McCarthy, J. F., Costa, T. J., and Liongosari, E. S. UniCast, Out-Cast & GroupCast: Three Steps Toward Ubiquitous, Peripheral Displays. In *Proc. UbiComp '01*, Springer (2001).
- [140] McCarthy, J. F., Farnham, S. D., Patel, Y., Ahuja, S., Norman, D., Hazlewood, W. R., and Lind, J. Supporting Community in Third Places with Situated Social Software. In *Proc. C&T '09*, ACM (2009).
- [141] McDonald, D. W., McCarthy, J. F., Soroczak, S., Nguyen, D. H., and Rashid, A. M. Proactive displays: Supporting awareness in fluid social environments. *TOCHI 14* (2008).
- [142] Memarovic, N. *Interacting Places - Networked Public Displays That Stimulate Community Interaction*. PhD thesis, Università della Svizzera Italiana, 2014.
- [143] Memarovic, N., Elhart, I., and Langheinrich, M. FunSquare: First Experiences with Autopoiesic Content. In *Proc. MUM '11*, ACM (2011).
- [144] Memarovic, N., Langheinrich, M., and Alt, F. The interacting places framework: conceptualizing public display applications that promote community interaction and place awareness. In *Proc. PerDis '12* (2012).

- [145] Memarovic, N., Langheinrich, M., Cheverst, K., Taylor, N., and Alt, F. P-LAYERS – A Layered Framework Addressing the Multifaceted Issues Facing Community-Supporting Public Display Deployments. *ACM Trans. Comput.-Hum. Interact* 20, 3 (2013).
- [146] Michelis, D., and Müller, J. The Audience Funnel: Observations of Gesture Based Interaction With Multiple Large Displays in a City Center. *Int. J. Human Computer Interaction* 27, 6 (2011).
- [147] Miyaoku, K., Higashino, S., and Tonomura, Y. C-blink: a hue-difference-based light signal marker for large screen interaction via any mobile terminal. In *Proc. UIST '04*, ACM (2004).
- [148] Moere, A. V., and Wouters, N. The Role of Context in Media Architecture. In *Proc. PerDis '12*, ACM (2012).
- [149] Müller, H., and Krüger, A. Learning Topologies of Situated Public Displays by Observing Implicit User Interactions. In *Universal Access in Human-Computer Interaction. Ambient Interaction*, C. Stephanidis, Ed., vol. 4555 of LNCS. Springer, 2007.
- [150] Müller, H. J. *Context adaptive digital signage in transitional spaces*. PhD thesis, University of Münster, 2008.
- [151] Müller, J., Alt, F., Michelis, D., and Schmidt, A. Requirements and design space for interactive public displays. In *Proc. MM '10*, ACM (2010).
- [152] Müller, J., Exeler, J., Buzeck, M., and Krüger, A. ReflectiveSigns: Digital Signs That Adapt to Audience Attention. In *Proc. Pervasive '09*, Springer (2009).
- [153] Müller, J., Jentsch, M., Kray, C., and Krüger, A. Exploring factors that influence the combined use of mobile devices and public displays for pedestrian navigation. In *Proc. NordiCHI '08*, ACM (2008).

-
- [154] Müller, J., and Krüger, A. Competing for your Attention: Negative Externalities in Digital Signage Advertising. In *Ambient Information Systems* (2007).
- [155] Müller, J., Krüger, A., and Kuflik, T. Maximizing the Utility of Situated Public Displays. In *User Modeling 2007*, C. Conati, K. McCoy, and G. Paliouras, Eds., vol. 4511 of *LNCS*. Springer, 2007.
- [156] Müller, J., Wilmsmann, D., Exeler, J., Buzeck, M., Schmidt, A., Jay, T., and Krüger, A. Display Blindness: The Effect of Expectations on Attention towards Digital Signage. In *Proc. Pervasive '09*, Springer (2009).
- [157] Munson, S. A., Rosengren, E., and Resnick, P. Thanks and tweets: comparing two public displays. In *Proc. CSCW '11*, ACM (2011).
- [158] Nakamura, T., Katayama, A., Yamamuro, M., and Sonehara, N. Fast watermark detection scheme for camera-equipped cellular phone. In *Proc. MUM '04*, ACM (2004).
- [159] Nakanishi, Y. Virtual prototyping using miniature model and visualization for interactive public displays. In *Proc. DIS '12*, ACM (2012).
- [160] Nancel, M., Wagner, J., Pietriga, E., Chapuis, O., and Mackay, W. Mid-air Pan-and-zoom on Wall-sized Displays. In *Proc. CHI '11*, ACM (2011).
- [161] Nissenbaum, H. *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford University Press, 2009.
- [162] Nittala, A. S., Yang, X.-D., Bateman, S., Sharlin, E., and Greenberg, S. Phonear: Interactions for mobile devices that hear

- high-frequency sound-encoded data. In *Proc. EICS '15*, ACM (2015).
- [163] O'Hara, K., Harper, R., Unger, A., Wilkes, J., Sharpe, B., and Jansen, M. TxtBoard: from text-to-person to text-to-home. In *Adj. Proc. CHI '05*, ACM (2005).
- [164] Ojala, T., Kostakos, V., Kukka, H., Heikkinen, T., Linden, T., Jurmu, M., Hosio, S., Kruger, F., and Zanni, D. Multipurpose Interactive Public Displays in the Wild: Three Years Later. *Computer* 45, 5 (2012).
- [165] Ojala, T., Kukka, H., Lindén, T., Heikkinen, T., Jurmu, M., Hosio, S., and Kruger, F. UBI-Hotspot 1.0: Large-Scale Long-Term Deployment of Interactive Public Displays in a City Center. In *Proc. ICIW '10*, IEEE (2010).
- [166] Ojala, T., Valkama, V., Kukka, H., Heikkinen, T., Lindén, T., Jurmu, M., Kruger, F., and Hosio, S. UBI-hotspots: sustainable ecosystem infrastructure for real world urban computing research and business. In *Proc. MEDES '10*, ACM (2010).
- [167] Olivier, P., Gilroy, S., Cao, H., Jackson, D., and Kray, C. Cross-modal Attention in Public-Private Displays. In *Proc. Pervasive Services '06*, IEEE (2006).
- [168] O'Neill, E., Woodgate, D., and Kostakos, V. Easing the Wait in the Emergency Room: Building a Theory of Public Information Systems. In *Proc. DIS '04*, ACM (2004).
- [169] Ostkamp, M., and Bauer, G. Multipleye – Concurrent Information Delivery on Public Displays. In *Adj. Proc. EuroITV '11*, COFAC (2011).
- [170] Ostkamp, M., Bauer, G., and Kray, C. Visual highlighting on public displays. In *Proc. PerDis '12*, ACM (2012).

- [171] Ostkamp, M., Heitmann, S., and Kray, C. Short-range optical interaction between smartphones and public displays. In *Proc. PerDis '15*, ACM (2015).
- [172] Ostkamp, M., Hülsermann, J., Kray, C., and Bauer, G. Using Mobile Devices to Enable Visual Multiplexing on Public Displays: Three Approaches Compared. In *Proc. MUM '13*, ACM (2013).
- [173] Ostkamp, M., and Kray, C. Prototyping mobile AR in immersive video environments. In *MobileHCI '13 Workshop*, ACM (2013).
- [174] Ostkamp, M., and Kray, C. Supporting Design, Prototyping, and Evaluation of Public Display Systems. In *Proc. EICS '14*, ACM (2014).
- [175] Ostkamp, M., Kray, C., and Bauer, G. Towards a Privacy Threat Model for Public Displays. In *Proc. EICS '15*, ACM (2015).
- [176] Ostkamp, M., Luzar, S., and Bauer, G. QR Codes on Curved Media Facades – Two Approaches For Inverse Distortion Based on Raytracing and Image Warping. In *Proc. GRAPP '14*, SCITEPRESS (2014).
- [177] Otto, O., Roberts, D., and Wolff, R. A review on effective closely-coupled collaboration using immersive CVE's. In *Proc. VRCA '06*, ACM (2006).
- [178] Palen, L., and Dourish, P. Unpacking "Privacy" for a Networked World. In *Proc. CHI '03*, ACM (2003).
- [179] Pearson, J., Robinson, S., and Jones, M. It's About Time: Smartwatches As Public Displays. In *Proc. CHI '15*, ACM (2015).
- [180] Peltonen, P., Kurvinen, E., Salovaara, A., Jacucci, G., Ilmonen, T., Evans, J., Oulasvirta, A., and Saarikko, P. It's Mine, Don't Touch!:

- interactions at a large multi-touch display in a city centre. In *Proc. CHI '08*, ACM (2008).
- [181] Penn, A., and Turner, A. Space syntax based agent simulation. In *Proc. Pedestrian and Evacuation Dynamics '01*, Open Access (2001).
- [182] Perry, M., Beckett, S., O'Hara, K., and Subramanian, S. WaveWindow: Public, Performative Gestural Interaction. In *Proc. ITS '10*, ACM (2010).
- [183] Perry, M., and O'Hara, K. Display-based activity in the workplace. In *Proc. INTERACT '03*, IOS Press (2003).
- [184] Raj, H., Gossweiler, R., and Milojevic, D. ContentCascade Incremental Content Exchange between Public Displays and Personal Devices. In *Proc. MobiQuitous '04*, IEEE (2004).
- [185] Röcker, C., Hinske, S., and Magerkurth, C. SPIROS - A system for privacy-enhanced information representation in smart home environments. In *Proc. IE '06*, vol. 1, IEEE (2006).
- [186] Röcker, C., Hinske, S., and Magerkurth, C. Intelligent Privacy Support for Large Public Displays. In *Universal Access in Human-Computer Interaction. Ambient Interaction*, C. Stephanidis, Ed., vol. 4555 of LNCS. Springer, 2007.
- [187] Rogers, A., David, E., Payne, T. R., and Jennings, N. R. An advanced bidding agent for advertisement selection on public displays. In *Proc. AAMAS '07*, ACM (2007).
- [188] Rogers, Y., and Brignull, H. Subtle ice-breaking: encouraging socializing and interaction around a large public display. In *CSCW '02 Workshop* (2002).

- [189] Rukzio, E., Müller, M., and Hardy, R. Design, implementation and evaluation of a novel public display for pedestrian navigation: the rotating compass. In *Proc. CHI '09*, ACM (2009).
- [190] Rukzio, E., Schmidt, A., Hussmann, H., and Informatics, M. An Analysis of the Usage of Mobile Phones for Personalized Interactions with Ubiquitous Public Displays. *Ubiquitous Public Displays, UBICOMP '04 Workshop* (2004).
- [191] Russell, D. M., and Gossweiler, R. On the Design of Personal & Communal Large Information Scale Appliances. In *Proc. UbiComp '01*, Springer (2001).
- [192] Sakurai, S., Kitamura, Y., Subramanian, S., and Kishino, F. A visibility control system for collaborative digital table. *Personal and Ubiquitous Computing* 13, 8 (2009).
- [193] Särkelä, H., Takatalo, J., May, P., Laakso, M., and Nyman, G. The movement patterns and the experiential components of virtual environments. *Int. J. Human-Computer Studies* 67, 9 (2009).
- [194] Schaad, A., and Borozdin, M. TAM2: Automated Threat Analysis. In *Proc. SAC '12*, ACM (2012).
- [195] Schaeffler, J. *Digital Signage: Software, Networks, Advertising, and Displays: A Primer for Understanding the Business*, 1 ed. Focal Press, 2008.
- [196] Schaub, F., Könings, B., Lang, P., Wiedersheim, B., Winkler, C., and Weber, M. PriCal: Context-adaptive Privacy in Ambient Calendar Displays. In *Proc. UbiComp '14*, ACM (2014).
- [197] Schellenbach, M., Krüger, A., Lövdén, M., and Lindenberger, U. A Laboratory Evaluation Framework for Pedestrian Navigation Devices. In *Proc. Mobility '07*, ACM (2007).

- [198] Schieferdecker, I., Grossmann, J., and Schneider, M. Model-Based Security Testing. In *MBT '12 Workshop*, EPTCS (2012).
- [199] Schmidt, A., Pfleging, B., Alt, F., Sahami, A., and Fitzpatrick, G. Interacting with 21st-Century Computers. *Pervasive Computing, IEEE 11*, 1 (2012).
- [200] Schneegass, S., and Alt, F. SenScreen: A Toolkit for Supporting Sensor-enabled Multi-Display Networks. In *Proc. PerDis '14*, ACM (2014).
- [201] Schneier, B. *Applied Cryptography: Protocols, Algorithms and Source Code in C*, 2 ed. John Wiley & Sons, 1995.
- [202] Schubert, F., Fath, T., and Haas, H. Coloured Video Code for In-Flight Data Transmission. In *Computer Vision Systems*, M. Chen, B. Leibe, and B. Neumann, Eds., vol. 7963 of *LNCS*. Springer, 2013.
- [203] Sharifi, M., Payne, T., and David, E. Public Display Advertising Based on Bluetooth Device Presence. In *MIRW, Human Computer Interaction with Mobile Devices and Services '06* (2006).
- [204] Shirazi, A. S., Winkler, C., and Schmidt, A. Flashlight interaction: a study on mobile phone interaction techniques with large displays. In *Proc. MobileHCI '09*, ACM (2009).
- [205] Shoemaker, G. B. D., and Inkpen, K. M. Single display privacyware: augmenting public displays with private information. In *Proc. CHI '01*, ACM (2001).
- [206] Silva, J., Ribeiro, Ó., Fernandes, J., Campos, J., and Harrison, M. The APEX Framework: Prototyping of Ubiquitous Environments Based on Petri Nets. In *Human-Centred Software Engineering*, R. Bernhaupt, P. Forbrig, J. Gulliksen, and M. Lárusdóttir, Eds., vol. 6409 of *LNCS*. Springer, 2010.

- [207] Silva, J. L., Campos, J., and Harrison, M. Formal analysis of ubiquitous computing environments through the apex framework. In *Proc. EICS '12*, ACM (2012).
- [208] Singh, P., Ha, H. N., Kuang, Z., Olivier, P., Kray, C., Blythe, P., and James, P. Immersive Video As a Rapid Prototyping and Evaluation Tool for Mobile and Ambient Applications. In *Proc. MobileHCI '06*, ACM (2006).
- [209] Singh, P., Ha, H. N., Olivier, P., Kray, C., Kuang, Z., Guo, A. W., Blythe, P., and James, P. Rapid prototyping and evaluation of intelligent environments using immersive video. In *MODIE, MobileHCI '06 Workshop* (2006).
- [210] Snowdon, C., and Kray, C. Exploring the Use of Landmarks for Mobile Navigation Support in Natural Environments. In *Proc. MobileHCI '09*, ACM (2009).
- [211] Snowdon, D., and Grasso, A. Diffusing Information in Organizational Settings: Learning from Experience. In *Proc. CHI '02*, ACM (2002).
- [212] Solove, D. J. A Taxonomy of Privacy. SSRN Scholarly Paper ID 667622, Social Science Research Network, Rochester, NY, 2005. <http://papers.ssrn.com/abstract=667622>, accessed: March 10, 2015.
- [213] Solove, D. J. 'I've Got Nothing to Hide' and Other Misunderstandings of Privacy. SSRN Scholarly Paper, Social Science Research Network, 2007. <http://papers.ssrn.com/abstract=998565>, accessed: September 25, 2014.
- [214] Stahl, C., and Hauptert, J. Simulating and evaluating public situated displays in virtual environment models. In *MODIE, MobileHCI '06 Workshop* (2006).

- [215] Steinberger, F., Foth, M., and Alt, F. Vote With Your Feet: Local Community Polling on Urban Screens. In *Proc. PerDis '14*, ACM (2014).
- [216] Stock, O., Zancanaro, M., Busetta, P., Callaway, C., Krüger, A., Kruppa, M., Kuflik, T., Not, E., and Rocchi, C. Adaptive, intelligent presentation of information for the museum visitor in PEACH. *User Modeling and User-Adapted Interaction* 17, 3 (2007).
- [217] Storz, O., Friday, A., Davies, N., Finney, J., Sas, C., and Sheridan, J. Public Ubiquitous Computing Systems: Lessons from the e-Campus Display Deployments. *Pervasive Computing, IEEE* 5, 3 (2006).
- [218] Sumi, Y., and Mase, K. AgentSalon: facilitating face-to-face knowledge exchange through conversations among personal agents. In *Proc. AGENTS '01*, ACM (2001).
- [219] Sutanto, J., Palme, E., Tan, C.-H., and Phang, C. W. Addressing the Personalization-privacy Paradox: An Empirical Assessment from a Field Experiment on Smartphone Users. *MIS Q.* 37, 4 (2013).
- [220] Taher, F., Cheverst, K., Harding, M., and Fitton, D. Formative Studies for Dynamic Wayfinding Support with In-building Situated Displays and Mobile Devices. In *Proc. MUM '09*, ACM (2009).
- [221] Taylor, N., and Cheverst, K. Social Interaction Around a Rural Community Photo Display. *Int. J. Human-Computer Studies* 67, 12 (2009).
- [222] Taylor, N., and Cheverst, K. Creating a rural community display with local engagement. In *Proc. DIS '11*, ACM (2010).

- [223] Taylor, N., and Cheverst, K. Supporting Community Awareness with Interactive Displays. *Computer* 45, 5 (2012).
- [224] Taylor, N., Cheverst, K., Fitton, D., Race, N. J. P., Rouncefield, M., and Graham, C. Probing Communities: Study of a Village Photo Display. In *Proc. OzCHI '07*, ACM (2007).
- [225] Ten Koppel, M., Bailly, G., Müller, J., and Walter, R. Chained displays: configurations of public displays can be used to influence actor-, audience-, and passer-by behavior. In *Proc. CHI '12*, ACM (2012).
- [226] The OWASP Foundation. OWASP Testing Guide, 2013. https://www.owasp.org/images/5/56/OWASP_Testing_Guide_v3.pdf, accessed: May 5, 2015.
- [227] The OWASP Foundation. OWASP Top 10 – 2013, 2013. <http://owasptop10.googlecode.com/files/OWASP%20Top%2010%20-%202013.pdf>, accessed: May 5, 2015.
- [228] Toch, E., Wang, Y., and Cranor, L. Personalization and privacy: a survey of privacy risks and remedies in personalization-based systems. *User Modeling and User-Adapted Interaction* 22, 1-2 (2012).
- [229] Unser, M. Sampling—50 Years After Shannon. *Proc. IEEE* 88, 4 (2000).
- [230] Valkanova, N., Walter, R., Vande Moere, A., and Müller, J. My-Position: Sparking Civic Discourse by a Public Interactive Poll Visualization. In *Proc. CSCW '14*, ACM (2014).
- [231] Vogel, D., and Balakrishnan, R. Interactive public ambient displays: transitioning from implicit to explicit, public to personal, interaction with multiple users. In *Proc. UIST '04*, ACM (2004).

- [232] Vogel, D., and Balakrishnan, R. Distant freehand pointing and clicking on very large, high resolution displays. In *Proc. UIST '05*, ACM (2005).
- [233] Wagner, D., Reitmayr, G., Mulloni, A., Drummond, T., and Schmalstieg, D. Real-Time Detection and Tracking for Augmented Reality on Mobile Phones. *IEEE Transactions on Visualization and Computer Graphics* 16, 3 (2010).
- [234] Walter, R., Bailly, G., and Müller, J. StrikeAPose: Revealing Mid-air Gestures on Public Displays. In *Proc. CHI '13*, ACM (2013).
- [235] Wang, M., Boring, S., and Greenberg, S. Proxemic peddler: a public advertising display that captures and preserves the attention of a passerby. In *Proc. PerDis '12*, ACM (2012).
- [236] Warren, S. D., and Brandeis, L. D. The Right to Privacy. *Harvard Law Review* 4, 5 (1890).
- [237] Weiser, M. The Computer for the 21st Century. *Scientific American* (1991).
- [238] Westin, A. F. *Privacy and freedom*. Atheneum, 1967.
- [239] Wilhelm, D., Fechner, T., Ostkamp, M., and Kray, C. Natural interaction with video environments using gestures and a mirror-image avatar. In *Proc. Interact '15*, Springer (2015).
- [240] Willis, K. D., Poupyrev, I., Hudson, S. E., and Mahler, M. SideBy-Side: Ad-hoc Multi-user Interaction with Handheld Projectors. In *Proc. UIST '11*, ACM (2011).
- [241] Wißner, M., Hammer, S., Kurdyukova, E., and André, E. Trust-based Decision-making for the Adaptation of Public Displays in Changing Social Contexts. *J. Trust Management* 1, 1 (2014).

- [242] Wolfe, J. M., and Horowitz, T. S. What attributes guide the deployment of visual attention and how do they do it? *Nat Rev Neurosci* 5, 6 (2004).
- [243] Xiao, R., Harrison, C., Willis, K. D., Poupyrev, I., and Hudson, S. E. Lumitrack: Low Cost, High Precision, High Speed Tracking with Projected M-sequences. In *Proc. UIST '13*, ACM (2013).
- [244] Xie, H., Filippidis, L., Gwynne, S., Galea, E. R., Blackshields, D., and Lawrence, P. J. Signage Legibility Distances as a Function of Observation Angle. *J. Fire Protection Engineering* 17, 1 (2007).
- [245] Yamada, T., and Kamitani, M. A method for detecting watermarks in print using smart phone: finding no mark. In *Proc. MoVid '13*, ACM (2013).
- [246] Youn, H., Park, C., and Lee, E. Security Based Survivability Risk Analysis with Extended HQPN. In *Proc. ICUIMC '11*, ACM (2011).
- [247] Zimmermann, A., Lorenz, A., and Oppermann, R. An Operational Definition of Context. In *Modeling and Using Context*, B. Kokinov, D. Richardson, T. Roth-Berghofer, and L. Vieu, Eds., vol. 4635 of *LNCS*. Springer, 2007.

Supplementary Material

Student Theses

The author supervised the following student theses:

[T1] Dominic Sondermann. Hervorheben von Informationen auf Großbildschirmen im öffentlichen Raum mithilfe von Augmented Reality. Fachhochschule Münster, 2012.

[T2] Malte Wesker. Effizienzvergleich verschiedener Interaktionsmodi für öffentliche Bildschirme. Fachhochschule Münster, 2013.

[T3] Sven Heitmann. Lichtblick — Optische Interaktion zwischen Mobiltelefon und Public Displays. Westfälische Wilhelms-Universität Münster, 2013.

[T4] Johanna Möllmann. Vergleich zweier Navigationsmethoden für immersive Video-Umgebungen. Westfälische Wilhelms-Universität Münster, 2014.

[T5] Jonas Hülsermann. Ein Time-Division-Multiplexing-Prototyp und seine Anwendung. Fachhochschule Münster, 2014.

[T6] Nicholas Schiestel. Umsetzung und Evaluierung einer Sprachsteuerung zur Navigation in einer IVE. Westfälische Wilhelms-Universität Münster, 2015.

Privacy Threat Model Relations

The following list contains all identified relations between the components of the privacy threat model for personalized public displays, see Subsection 10.1.2, pp. 190. Each relation is structured as follows: ($\langle \text{agent 1} \rangle, \langle \text{agent 2} \rangle, \dots$), ($\langle \text{threat 1} \rangle, \langle \text{threat 2} \rangle, \dots$), ($\langle \text{weakness 1} \rangle, \langle \text{weakness 2} \rangle, \dots$), ($\langle \text{effect 1} \rangle, \langle \text{effect 2} \rangle, \dots$), ($\langle \text{countermeasure 1} \rangle, \langle \text{countermeasure 2} \rangle, \dots$). Sets of similar components are marked by brackets (e.g., $\langle \text{A4} \rangle, \langle \text{A5} \rangle$), while brackets around sets containing only one element are optional (e.g., $\langle \text{A4} \rangle = \langle \text{A4} \rangle$). Due to the chosen implementation of the prototype (see Section 11.1), the threats had to be subdivided according to the agents: The symbol $\langle \text{T0.1} \rangle$ stands for the threat “normal usage” that can only be applied by agent No. 1, i.e., “Alice, Bob, ...,” see Table 10.9. Thus, the number behind the decimal point may be neglected when interpreting an individual relation.

- ($\langle \text{A1} \rangle, \langle \text{T0.1} \rangle, \{ \langle \text{W1} \rangle, \langle \text{W2} \rangle, \langle \text{W3} \rangle, \langle \text{W10} \rangle, \langle \text{W11} \rangle \}, \{ \langle \text{E1} \rangle, \langle \text{E2} \rangle, \langle \text{E3} \rangle, \langle \text{E4} \rangle, \langle \text{E5} \rangle, \langle \text{E6} \rangle \}, \{ \langle \text{C1} \rangle, \langle \text{C14} \rangle, \langle \text{C15} \rangle, \langle \text{C3} \rangle, \langle \text{C10} \rangle, \langle \text{C19} \rangle, \langle \text{C23} \rangle, \langle \text{C22} \rangle, \langle \text{C10} \rangle, \langle \text{C14} \rangle, \langle \text{C18} \rangle, \langle \text{C21} \rangle, \langle \text{C7} \rangle, \langle \text{C14} \rangle, \langle \text{C21} \rangle, \langle \text{C6} \rangle, \langle \text{C2} \rangle, \langle \text{C3} \rangle, \langle \text{C4} \rangle, \langle \text{C5} \rangle, \langle \text{C5} \rangle, \langle \text{C9} \rangle, \langle \text{C14} \rangle, \langle \text{C24} \rangle, \langle \text{C4} \rangle, \langle \text{C5} \rangle, \langle \text{C15} \rangle, \langle \text{C21} \rangle, \langle \text{C6} \rangle, \langle \text{C14} \rangle, \langle \text{C21} \rangle$)
- ($\langle \text{A3} \rangle, \langle \text{T0.2} \rangle, \langle \text{W2} \rangle, \{ \langle \text{E1} \rangle, \langle \text{E6} \rangle \}, \{ \langle \text{C1} \rangle, \langle \text{C23} \rangle, \langle \text{C3} \rangle, \langle \text{C10} \rangle, \langle \text{C19} \rangle, \langle \text{C23} \rangle, \langle \text{C2} \rangle, \langle \text{C8} \rangle, \langle \text{C10} \rangle, \langle \text{C12} \rangle, \langle \text{C15} \rangle, \langle \text{C25} \rangle, \langle \text{C15} \rangle, \langle \text{C20} \rangle, \langle \text{C24} \rangle, \langle \text{C17} \rangle, \langle \text{C19} \rangle, \langle \text{C13} \rangle, \langle \text{C22} \rangle, \langle \text{C2} \rangle, \langle \text{C3} \rangle, \langle \text{C4} \rangle, \langle \text{C5} \rangle, \langle \text{C5} \rangle, \langle \text{C9} \rangle, \langle \text{C14} \rangle, \langle \text{C24} \rangle, \langle \text{C4} \rangle, \langle \text{C5} \rangle, \langle \text{C15} \rangle, \langle \text{C21} \rangle, \langle \text{C6} \rangle, \langle \text{C14} \rangle, \langle \text{C21} \rangle$)
- ($\{ \langle \text{A4} \rangle, \langle \text{A5} \rangle \}, \langle \text{T0.3} \rangle, \langle \text{W2} \rangle, \{ \langle \text{E1} \rangle, \langle \text{E6} \rangle \}, \{ \langle \text{C1} \rangle, \langle \text{C23} \rangle, \langle \text{C3} \rangle, \langle \text{C10} \rangle, \langle \text{C19} \rangle, \langle \text{C23} \rangle, \langle \text{C2} \rangle, \langle \text{C8} \rangle, \langle \text{C10} \rangle, \langle \text{C12} \rangle, \langle \text{C15} \rangle, \langle \text{C25} \rangle, \langle \text{C15} \rangle, \langle \text{C20} \rangle, \langle \text{C24} \rangle, \langle \text{C17} \rangle, \langle \text{C19} \rangle, \langle \text{C13} \rangle, \langle \text{C22} \rangle, \langle \text{C2} \rangle, \langle \text{C3} \rangle$)

- <C4>, <C5>, <C5>, <C9>, <C14>, <C24>, <C4>, <C5>, <C15>, <C21>, <C6>, <C14>, <C21>))
- (<A2>, <T1.1>, <W1>, {<E1>, <E4>, <E5>, <E6>}, {<C1>, <C8>, <C12>, <C24>, <C20>, <C24>, <C17>, <C19>, <C12>, <C20>, <C24>, <C19>, <C6>, <C23>, <C14>, <C15>})
 - (<A2>, <T1.2>, <W2>, {<E2>, <E3>, <E4>, <E5>, <E6>}, {<C1>, <C8>, <C12>, <C24>, <C20>, <C24>, <C17>, <C19>, <C12>, <C20>, <C24>, <C19>, <C6>, <C12>, <C22>, <C10>, <C18>})
 - (<A2>, <T1.3>, <W3>, {<E2>, <E4>, <E5>, <E6>}, {<C1>, <C8>, <C12>, <C12>, <C22>})
 - (<A2>, <T1.4>, <W4>, {<E1>, <E3>, <E5>, <E6>}, {<C1>, <C24>, <C20>, <C24>, <C17>, <C19>, <C12>, <C22>, <C17>, <C19>, <C20>, <C24>, <C17>, <C19>, <C6>})
 - (<A2>, <T1.5>, <W5>, {<E1>, <E3>, <E5>, <E6>}, {<C1>, <C17>, <C23>, <C15>, <C20>, <C24>, <C17>, <C19>, <C6>, <C2>, <C3>, <C4>, <C5>, <C5>, <C9>, <C14>, <C24>, <C4>, <C5>, <C15>, <C21>})
 - (<A2>, <T1.6>, <W6>, {<E1>, <E5>, <E6>}, {<C1>, <C24>})
 - (<A2>, <T1.7>, <W11>, {<E1>, <E2>, <E4>, <E5>, <E6>}, {<C1>, <C24>, <C20>, <C24>, <C12>, <C22>, <C17>, <C19>, <C20>, <C24>, <C15>, <C20>, <C24>, <C17>, <C19>, <C6>})
 - (<A2>, <T2.1>, <W1>, {<E1>, <E2>, <E3>, <E4>, <E5>, <E6>}, {<C1>, <C8>, <C12>, <C24>, <C20>, <C24>, <C17>, <C19>, <C12>, <C20>, <C24>, <C19>, <C6>})
 - (<A2>, <T2.2>, <W2>, {<E2>, <E3>, <E4>, <E5>, <E6>}, {<C1>, <C8>, <C12>, <C24>, <C20>, <C24>, <C17>, <C19>, <C12>, <C20>, <C24>, <C19>, <C6>})

- (<A2>, <T2.3>, <W3>, {<E2>, <E6>}, {<C1>, <C8>, <C12>})
- (<A2>, <T2.4>, <W4>, {<E1>, <E4>, <E5>, <E6>}, {<C1>, <C17>, <C19>, <C20>, <C24>})
- (<A2>, <T2.5>, <W5>, {<E1>, <E3>, <E5>, <E6>}, {<C1>, <C15>, <C20>, <C24>})
- (<A2>, <T2.6>, <W6>, {<E1>, <E4>, <E5>, <E6>}, {<C1>, <C24>})
- (<A2>, <T2.7>, <W7>, {<E1>, <E2>, <E3>, <E4>, <E5>, <E6>}, {<C1>, <C24>, <C20>, <C24>, <C17>, <C19>, <C17>, <C19>, <C20>, <C24>})
- (<A2>, <T2.8>, <W9>, {<E1>, <E2>, <E3>, <E4>, <E5>, <E6>}, {<C1>, <C3>, <C10>, <C19>, <C23>, <C10>, <C18>, <C2>, <C8>, <C10>, <C12>, <C15>, <C25>, <C5>, <C9>, <C14>, <C24>})
- (<A2>, <T2.9>, <W11>, {<E1>, <E5>, <E6>}, {<C1>, <C20>, <C24>, <C17>, <C23>, <C3>, <C10>, <C19>, <C23>, <C12>, <C22>, <C17>, <C19>, <C6>, <C6>, <C14>, <C21>})
- ({<A3>, <A5>}, <T2.10>, <W1>, {<E1>, <E2>, <E3>, <E4>, <E5>, <E6>}, {<C1>, <C8>, <C12>, <C24>, <C20>, <C24>, <C17>, <C19>, <C12>, <C20>, <C24>, <C19>, <C6>})
- ({<A3>, <A5>}, <T2.11>, <W2>, {<E1>, <E2>, <E3>, <E4>, <E5>, <E6>}, {<C1>, <C8>, <C12>, <C24>, <C20>, <C24>, <C17>, <C19>, <C12>, <C20>, <C24>, <C19>, <C6>})
- ({<A3>, <A5>}, <T2.12>, <W4>, {<E1>, <E2>, <E3>, <E4>, <E5>, <E6>}, {<C1>, <C17>, <C19>, <C20>, <C24>})
- (<A1>, <T3.1>, {<W8>}, {<E2>, <E4>, <E5>, <E6>}, {<C1>, }, (<A1>, <T3.1>, {<W8>, <W9>, <W10>}, {<E2>, <E4>, <E5>, <E6>}, {<C1>, <C8>, <C12>, <C14>, <C15>, <C12>, <C22>, <C10>, <C18>, <C2>, <C8>, <C10>, <C12>, <C15>, <C25>, <C17>, <C19>, <C7>, <C14>,

- <C21>, <C5>, <C9>, <C14>, <C24>, <C4>, <C5>, <C15>, <C21>, <C6>, <C14>, <C21>))
- (<A2>, <T3.2>, <W7>, {<E1>, <E2>, <E3>, <E4>, <E5>, <E6>}, {<C1>, <C20>, <C24>})
 - (<A1>, <T4.1>, {<W1>, <W2>}, {<E1>, <E5>, <E6>}, {<C1>, <C3>, <C10>, <C19>, <C23>}), (<A1>, <T4.1>, {<W10>}, {<E1>, <E5>, <E6>}, {<C1>, <C4>, <C5>, <C15>, <C21>, <C6>, <C14>, <C21>})
 - (<A2>, <T4.2>, {<W1>, <W2>}, {<E1>, <E5>, <E6>}, {<C3>, <C10>, <C19>, <C23>})
 - (<A2>, <T4.3>, {<W3>}, {<E1>, <E5>, <E6>}, {<C1>, <C6>, <C2>, <C3>, <C4>, <C5>})
 - (<A2>, <T4.4>, {<W4>}, {<E1>, <E2>, <E3>, <E4>, <E5>, <E6>}, {<C1>, <C13>}), (<A2>, <T4.4>, {<W4>, <W5>, <W11>}, {<E1>, <E2>, <E3>, <E4>, <E5>, <E6>}, {<C1>, <C17>, <C17>, <C19>, <C6>})
 - (<A2>, <T4.5>, {<W6>}, {<E1>, <E5>, <E6>}, {<C1>, <C24>, <C20>, <C24>, <C17>, <C12>, <C20>, <C24>, <C3>, <C10>, <C19>, <C23>, <C10>, <C18>, <C2>, <C8>, <C10>, <C12>, <C15>, <C25>, <C20>, <C24>, <C17>, <C19>, <C7>, <C14>, <C21>, <C6>})
 - (<A2>, <T4.6>, {<W11>}, {<E1>, <E5>, <E6>}, {<C1>, <C17>, <C3>, <C10>, <C19>, <C23>, <C13>, <C6>})
 - ({<A1>, <A2>, <A3>, <A7>, <T4.7>, {<W9>}, {<E1>, <E5>, <E6>}, {<C1>, <C5>, <C9>, <C14>, <C24>, <C4>, <C5>, <C15>, <C21>})
 - ({<A1>, <A3>, <A4>, <A5>, <T5.1>, <W1>, {<E1>, <E2>, <E5>, <E6>}, {<C1>, <C17>, <C14>, <C15>, <C3>, <C10>, <C19>, <C23>, <C17>, <C19>, <C13>, <C22>})

- (<A7>, <T5.7>, <W3>, {<E2>, <E5>, <E6>}, {<C1>, <C22>, <C10>, <C14>, <C18>, <C21>, <C7>, <C14>, <C21>})
- ({<A1>, <A2>, <A3>, <A7>}, <T5.8>, {<W9>}, {<E1>, <E5>, <E6>}, {<C1>, <C5>, <C9>, <C14>, <C24>})
- (<A1>, <T6.1>, {<W1>}, {<E1>, <E2>, <E5>, <E6>}, {<C1>, <C23>, <C14>, <C15>}), (<A1>, <T6.1>, {<W9>, <W10>}, {<E1>, <E2>, <E5>, <E6>}, {<C1>, <C12>, <C20>, <C24>, <C22>, <C10>, <C14>, <C18>, <C21>, <C7>, <C14>, <C21>, <C6>, <C6>, <C14>, <C21>})
- ({<A2>, <A3>}, <T6.2>, {<W1>}, {<E1>, <E2>, <E3>, <E4>, <E5>, <E6>}, {<C1>, <C23>, <C14>, <C15>}), ({<A2>, <A3>}, <T6.2>, {<W2>}, {<E1>, <E2>, <E3>, <E4>, <E5>, <E6>}, {<C1>, <C2>, <C8>, <C10>, <C12>, <C15>, <C25>})
- (<A2>, <T6.3>, {<W4>}, {<E1>, <E2>, <E3>, <E4>, <E5>, <E6>}, {<C1>, <C17>}), (<A2>, <T6.3>, {<W5>}, {<E1>, <E2>, <E3>, <E4>, <E5>, <E6>}, {<C1>, <C17>}), (<A2>, <T6.3>, {<W11>}, {<E1>, <E2>, <E3>, <E4>, <E5>, <E6>}, {<C1>, <C24>, <C17>, <C23>, <C3>, <C10>, <C19>, <C23>, <C15>, <C20>, <C24>, <C17>, <C19>, <C7>, <C14>, <C21>, <C6>, <C2>, <C3>, <C4>, <C5>, <C4>, <C5>, <C15>, <C21>})
- (<A2>, <T7.1>, {<W1>, <W2>}, {<E1>, <E2>, <E3>, <E5>, <E6>}, <C1>)
- ({<A2>, <A3>}, <T7.2>, {<W4>}, {<E1>, <E2>, <E3>, <E5>, <E6>}, {<C1>, <C17>, <C19>}), ({<A2>, <A3>}, <T7.2>, {<W5>}, {<E1>, <E2>, <E3>, <E5>, <E6>}, {<C1>, <C17>})
- (<A2>, <T7.3>, <W9>, {<E1>, <E2>, <E3>, <E5>, <E6>}, {<C1>, <C4>, <C5>, <C15>, <C21>})
- ({<A4>, <A5>, <A6>}, <T7.4>, {<W9>, <W10>}, {<E1>, <E2>, <E3>, <E5>, <E6>}, {<C1>, <C4>, <C5>, <C15>, <C21>})

- ({<A1>, <A2>, <A3>, <A4>, <A5>, <A6>, <A7>}, {<T0.1>, <T0.2>, <T0.3>, <T1.1>, <T1.2>, <T1.3>, <T1.4>, <T1.5>, <T1.6>, <T1.7>, <T2.1>, <T2.2>, <T2.3>, <T2.4>, <T2.5>, <T2.6>, <T2.7>, <T2.8>, <T2.9>, <T2.10>, <T2.11>, <T2.12>, <T3.1>, <T3.2>, <T4.1>, <T4.2>, <T4.3>, <T4.4>, <T4.5>, <T4.6>, <T4.7>, <T5.1>, <T5.2>, <T5.3>, <T5.4>, <T5.5>, <T5.6>, <T5.7>, <T5.8>, <T6.1>, <T6.2>, <T6.3>, <T7.1>, <T7.2>, <T7.3>, <T7.4>}, {<W1>, <W2>, <W3>, <W4>, <W5>, <W8>, <W10>, <W11>}, {<E1>, <E2>, <E3>, <E4>, <E5>, <E6>}, {<C11>})
- ({<A1>, <A2>, <A3>, <A4>, <A5>, <A6>, <A7>}, {<T0.1>, <T0.2>, <T0.3>, <T1.1>, <T1.2>, <T1.3>, <T1.4>, <T1.5>, <T1.6>, <T1.7>, <T2.1>, <T2.2>, <T2.3>, <T2.4>, <T2.5>, <T2.6>, <T2.7>, <T2.8>, <T2.9>, <T2.10>, <T2.11>, <T2.12>, <T3.1>, <T3.2>, <T4.1>, <T4.2>, <T4.3>, <T4.4>, <T4.5>, <T4.6>, <T4.7>, <T5.1>, <T5.2>, <T5.3>, <T5.4>, <T5.5>, <T5.6>, <T5.7>, <T5.8>, <T6.1>, <T6.2>, <T6.3>, <T7.1>, <T7.2>, <T7.3>, <T7.4>}, {<W1>, <W6>, <W9>, <W11>}, {<E1>, <E2>, <E3>, <E5>, <E6>}, {<C16>})

Cover Letter for the Qualitative Evaluation of CI

Dear ...,

the Situated Computing Lab at Münster University is working on an Interactive Public Display Privacy Threat Model. This work is part of Morin Ostkamp's PhD. Morin is interested in designing "privacy-preserving personalized public display systems". One of his research questions is: "How to support the design process of public displays to foster privacy-preserving systems?" His vision is to develop a holistic toolkit, that is based on a public display model, a threat model, and a set of privacy-preserving countermeasures. Designers of public display systems may integrate that toolkit into their workflow to design, prototype, and evaluate public display systems, that do not pose a threat to the user's privacy.

We have built a web application that implements a first version of the threat model. We recognize you as an expert in the domain of public displays and privacy and are thus interested in your opinion of our work. We kindly ask you to spend approximately 10 minutes to go to <http://ipdptm.se-labor.de>, experience our prototype, and answer these three questions afterwards:

1. What do you think about the extend of the threat model? Is it comprehensive or are we missing something important, e.g., a particular agent or a particular threat?
2. How can the threat model help designers of public displays to design privacy-preserving systems?
3. What are your comments or remarks that you would like to share with us in regards to the (interactive) threat model?

We highly appreciate your feedback and are looking forward to hearing from you.

Kind regards,

...

Curriculum Vitae

Designing Privacy-Preserving Personalized Public Display Systems

Morin Ostkamp

Digital public displays are a popular means of communication nowadays. They are commonly used as information outlets at traffic hubs, shopping malls, or public places in general. They feature some key advantages in comparison to other media types. For example, they are more flexible and more up-to-date than paper based approaches, and they usually have bigger screens than personal devices. Additionally, they are visually prominent, provide broad accessibility, and are situated in a certain context.

Clearly, showing users content that is “relevant” to them is an important issue. For example, due to a lack of relevant content, many people have developed a blindness towards public displays. Personal content is often regarded as relevant, but that calls for certain means of privacy in turn. This thesis focuses on designing “privacy-preserving personalized public display systems.” It addresses three research questions: What are main privacy threats on public displays? What are countermeasures to those privacy threats? How to support the design process of public displays?

Three tangible contributions address each research question: a privacy threat model, a list and classification of existing countermeasures accompanied by three novel countermeasures, and an integrated process. The latter provides a new methodology to design, prototype, and evaluate privacy-preserving personalized public display systems. Designers and researchers can use these contributions to create public displays, that do not pose a threat to the user's privacy.