Next Generation Analytics for Open Pervasive Display Networks

Mateusz Andrzej Mikusz

School of Computing and Communications
Lancaster University

This dissertation is submitted for the degree of
Doctor of Philosophy

October 2018

Declaration

This thesis has been written by myself and has not been submitted in support of an application for another degree at this or any other university. Excerpts of this dissertation have been published in the following conference manuscripts and academic publications on which I was a leading or contributing author (in reverse chronological order): [Mik+18a; Mik+18b; Mik+18c; Mik+18d; MCD17; Elh+17; Mik+16; Mik+15; MCD15; Cli+14].

The research conducted in the context of this thesis was partially funded through the Future and Emerging Technologies (FET) programme within the 7th Framework Programme for Research of the European Commission, under FET grant number: 612933 (RECALL), by the UK EPSRC under grant numbers EP/N023234/1 (PETRAS) and EP/N028228/1 (PACTMAN), and by a Google Faculty Research Award.

Some portions of this thesis have been developed in collaboration with colleagues from Lancaster University and other institutions. The e-Campus display network was subject to prior research described in [FDE12; Cli+13; Dav+14]. In particular, the Yarely signage player was designed and developed by Clinch et al. [Cli+13]. The original system architecture and design of Tacita was first introduced by Davies et al. [Dav+14] and reported in detail in [Cli13]. The Tacita mobile client described in Section 3.4.1.3 (p. 69) was implemented by Peter Shaw. The following personalisable applications introduced in Table 6.5 (p. 144) were implemented by Peter Shaw and Ludwig Trotter: Bus Departures, World Clock, News, Live TV, Twitter News Feed, and Pictures. The walk-by evaluations of the Bluetooth Low Energy beacon performance described in Section 6.3.2.1 (p. 134) were conducted with the help of Peter Shaw, Ludwig Trotter and Petteri Nurmi. The prototype mobile application to capture beacon detection latencies described in Section 6.3.2.2 (p. 135) was implemented by Peter Shaw. The map provider and mapping schema of public displays (Section 3.4.1.2, p. 69) was designed in collaboration with Nigel Davies, Peter Shaw, Ivan Elhart and Marc Langheinrich. In the context of Tacita, my contributions were focussed on designing the revised overall system architecture and associated application programming interfaces spanning across the mobile phone client, trusted content providers, display gateway and display nodes. The Wi-Fi fingerprinting approach to tracking individuals as part of the infrastructure-based tracking approach (Sections 3.4.2, p. 73; and 6.4, p. 149) was provided by LiveLabs [Jay+16].

Mateusz Andrzej Mikusz
October 2018

Acknowledgements

First and foremost, I would like to express my deepest gratitude to my supervisor Professor Nigel Davies. His continuous guidance and mentoring throughout my Ph.D., abilities of always having a clear vision of the research and often seeing the unobvious were of enormous help and truly inspiring. The endless coffee and late-night paper writing sessions were highly appreciated! I also owe my deepest thanks for the opportunities that were given beyond the scope of 'just' conducting my Ph.D.: collaborations with other research institutions, spending time abroad and contributing to a wide range of projects. This ultimately allowed me to gain an enormous amount of experience and to develop and foster my future career. I would like thank Professor Rajesh Krishna Balan and Professor Youngki Lee who very kindly hosted me as a visiting researcher at LiveLabs at Singapore Management University, and invited me for a subsequent visit that allowed me to continue working on our collaborative projects. I am extending my gratitude to my previous supervisors Professor Albrecht Schmidt and Professor Florian Alt for introducing me to Lancaster, and giving me the cherished opportunity to conduct a semester abroad in the context of my undergraduate studies that eventually led to my Ph.D.

One of my main reasons for conducting my Ph.D. at Lancaster (despite the good weather, of course) was the amazing and unique work environment at the School of Computing and Communications and the Data Science Institute. Special thanks is given to Professor Adrian Friday and Dr. Sarah Clinch who (especially early on) provided me with assistance to get a head start on e-Campus. Just like any other system, a display network would be worthless without an active user base. I am grateful to all e-Campus users with a special mention of the Internal Communications team who helped us turn the initially small deployment into the dedicated signage network of the University. Of course, one of the major pillars of an amazing work environment are the colleagues with whom I have shared an office space. Therefore, special thanks to all current and previous members of 'Team D23' including Mike Harding who always provided help and guidance to everyone, Peter Shaw and Sarah Clinch with whom I shared endless hours debugging e-Campus components and worked hard to bring e-Campus forward, and Ludwig Trotter for helping to recruit lots of Tacita users.

Throughout my Ph.D., I have worked on a number of national and international projects and collaborated with many other researchers. In particular, I would like to extend my thankfulness to all project members of PD-NET, RECALL, PETRAS and PACTMAN. Working on these projects was rewarding and very insightful and allowed me to foster many still ongoing collaborations. Special thanks to Professor Marc Langheinrich and Dr. Ivan Elhart for the efforts of making Tacita work across different deployment sites.

The completion of my studies would have been impossible without permanent support and encouragement of my close family. I would like to thank Katharina for her tireless support, motivation and also her patience during countless late-night working hours. I would also like to express my deepest gratitude to my parents Bogumila and Tomasz, my brother Marcin and his wife Ewelina and their family, who provided me with endless opportunities, continuous support and encouragement throughout my entire educational career.

Abstract

Next Generation Analytics for Open Pervasive Display Networks
Mateusz Andrzej Mikusz

Dipl.-Inf. (University of Stuttgart, 2013)
Submitted for the degree of Doctor of Philosophy.

October 2018

Public displays and digital signs are becoming increasingly widely deployed as many spaces move towards becoming highly interactive and augmented environments. Market trends suggest further significant increases in the number of digital signs and both researchers and commercial entities are working on designing and developing novel uses for this technology. Given the level of investment, it is increasingly important to be able to understand the effectiveness of public displays. Current state-of-the-art analytics technology is limited in the extent to which it addresses the challenges that arise from display deployments becoming *open* (increasing numbers of stakeholders), *networked* (viewer engagement across devices and locations) and *pervasive* (high density of displays and sensing technology leading to potential privacy threats for viewers).

In this thesis, we provide the first exploration into achieving next generation display analytics in the context of open pervasive display networks. In particular, we investigated three areas of challenge: *analytics data capture*, *reporting* and *automated use of analytics data*. Drawing on the increasing number of stakeholders, we conducted an extensive review of related work to identify data that can be captured by individual stakeholders of a display network, and highlighted the opportunities for gaining insights by combining datasets owned by different stakeholders. Additionally, we identified the importance of *viewer-centric analytics* that use traditional display-oriented analytics data combined with viewer mobility patterns to produce entirely new sets of analytics reports. We explored a range of approaches to generating viewer-centric analytics including the use of mobility models as a way to create 'synthetic analytics' – an approach that provides highly detailed analytics whilst preserving viewer privacy.

We created a collection of novel viewer-centric analytics reports providing insights into how viewers experience a large network of pervasive displays including reports regarding

the effectiveness of displays, the visibility of content across the display network, and the visibility of content to viewers. We further identified additional reports specific to those display networks that support the delivery of personalised content to viewers. Additionally, we highlighted the similarities between digital signage and Web analytics and introduced novel forms of digital signage analytics reports created by leveraging existing Web analytics engines.

Whilst the majority of analytics systems focus solely on the capture and reporting of analytics insights, we additionally explored the automated use of analytics data. One of the challenges in open pervasive display networks is accommodating potentially competing content scheduling constraints and requirements that originate from the large number of stakeholders – in addition to contextual changes that may originate from analytics insights. To address these challenges, we designed and developed the first lottery scheduling approach for digital signage providing a means to accommodate potentially conflicting scheduling constraints, and supporting context- and event-based scheduling based on analytics data fed back into the digital sign.

In order to evaluate the set of systems and approaches presented in this thesis, we conducted large-scale, long-term trials allowing us to show both the technical feasibility of the systems developed and provide insights into the accuracy and performance of different analytics capture technologies. Our work provides a set of tools and techniques for next generation digital signage analytics and lays the foundation for more general people-centric analytics that go beyond the domain of digital signs and enable unique analytical insights and understanding into how users interact across the physical and digital world.

Table of Contents

List of Figures

List of Tables

List of Listings

Nomenclature

Acronyms / Abbreviations

ACM	Association for Computing Machinery
AIM	Anonymous Impression Metrics
API	Application Programming Interface
AVA	Anonymous Video Analytics
BLE	Bluetooth Low Energy
CDS	Content Descriptor Set
IoT	Internet of Things
QR	Quick Response (Code)
REST	Representational State Transfer
RFID	Radio-frequency Identification
RTLS	Real-time Location Tracking
TID	Tracking Identifier
UMP	Universal Measurement Protocol
URL	Uniform Resource Locator
UUID	Universally Unique Identifier
XML	Extensible Markup Language

Chapter 1

Introduction

1.1 The Emergence of Pervasive Display Networks

Market reports predict a dramatic increase in the number of digital signs and public displays deployed in urban spaces [Mar17] with a total of "38 million connected digital screens in use worldwide" [Mar17]. By 2021, the number of deployed digital signs is expected to have more than doubled to a total of around 87 million digital signs [Dig17; Mar17]. In line with the increase in the number of deployed digital signs, the market worth for public displays is expected to grow to up to $23 billion indicating the increasing importance of digital signs [Glo17] – despite the high popularity of mobile phones as a way to access information. In addition to the growth in the number of deployments, we also observe a growth in the size of individual digital signs. For example, large display walls such as the high-definition display wall at Suntec Singapore (see Figure 1.1) are likely to lead to a world in which digital signage will become omnipresent to passers-by and integrated into their daily life [Mik+18a].

Digital signs and public displays are increasingly often also offering a high level of interactivity to the viewer. Early work in this space conducted by Vogel and Balakrishnan provided insights into interaction models that digital signs could support in future: digital signs automatically adapt the displayed content based on the spatial location of the viewer in relation to the digital sign [VB04]. Such interaction models have been further developed to, for example,

Figure 1.1: World's largest high-definition display wall at the Suntec Singapore convention centre.

additionally consider viewer identity to enable the delivery of personalised content [Gre+11]. Novel interaction techniques for public displays are constantly developed and include gaze, gesture, forms of remote control and direct touch input [She+14] – highlighting the trend towards highly interactive public displays. Interactive public displays can be often found situated in railway stations, airports and shopping malls in the form of kiosks allowing the passers-by to search for a specific piece of information.

Besides the increase in the scale and size of digital signs and the development of novel interaction modalities, researchers envision new directions in which digital signage will develop – changing the nature and characteristics of pervasive display deployments. Davies et al. [Dav+12] introduced the vision of *open display networks* as the "new communications medium for the 21st century" [Dav+12] in which a large number of distinct stakeholders contribute content, infrastructure and devices to a large, common displays network. Davies et al. envision future display deployments that feature similar characteristics as modern mobile phone ecosystems in which open platforms such as application stores exist that allow third-parties including developers, content creators, manufacturers and users to contribute equally to the market [Dav+12]. Such systems will need to feature open application programming interfaces that can be used by third-party developers to create applications without detailed knowledge about the specifics of certain display deployments and write applications that can be supported across large display deployments as part of an open network. Currently, commercial signage networks are closed and cannot be easily accessed by external developers and third-parties. Moving toward an open display network would substantially increase the potential set of stakeholders and contributors. It is argued that an open application platform for displays will lead to a comparable dramatic increase in the number of applications and content that has become available to mobile phones through the application store models [Dav+12]. The *networked* characteristic of future display deployments is particularly relevant to describe the potential scale and dimensions of display deployments especially with the current expectations of display deployments to grow to close to hundred million devices [Dig17; Mar17; Dav+12] – an open displays platform and application store for public displays would immediately provide access to a large number of displays to third-party developers, content creators and display owners [Cli+14]. Davies et al. point out that "openness does not stop at developers: public display systems should also be open to content from users" [Dav+12]. With displays becoming ubiquitously integrated into the daily life of users – both in public and private spaces – the contribution of users to the display network and content displayed will become important to ensure displays are useful to the user and provide a clear benefit through display personalisation [Dav+14].

To illustrate the future use of public displays in the context of such future networks, Davies, Clinch, and Alt [DCA14] have developed a number of scenarios. The following scenario is an example of the use of 'open displays' in the context of a small independent retail shop.

> A local fruit shop has a large number of fresh strawberries left to sell while the shop is due to close in a few hours already. The owner of the shop knows that

> strawberries are currently very popular and sell well at the competition, but has to advertise his strawberries in the local community. He quickly creates a short advertisement and submits it to the "local shopper incentive" scheme in which viewers who see the advert can retrieve a discount voucher for the local fruit shop. The discount voucher is also used to track the number of purchases retrieved from the advert and the shop owner has to pay a fee based on the number of new customers who used the discount. [DCA14]

The scenario above highlights the contrast to large commercial advertisement networks that typically feature more complex infrastructure and accessibility models – making it challenging for the local shop owner to quickly access display resources. It is further envisioned that future display networks will aim to influence viewer behaviour by showing more appropriate and personalised content for the viewer, as described in the following scenario:

> Jack is participating in a local walk-to-school programme that aims to increase the fitness of elementary school children. To encourage participation, the a game has been deployed on the local public display network. On his walk to school, Jack walks by a number of displays that provide him with a cartoon character and update on his fitness progress. By walking by a display, Jack also collects points on his mobile phone that can be exchanged for a cartoon book. [Dav+12]

The scenario above illustrates the opportunities of providing personalised content to passers-by and supporting future forms of user interaction and engagement modalities [Dav+12]. Display networks provide a platform to schedule content for individuals regardless of the displays' schedules but with respect to the (constantly changing) context of the display [DCA14]. The interactions between the passers-by and the display application enables the stakeholders of the display network to run novel forms of campaigns that ultimately lead to a behavioural change of the viewer [DCA14]. Providing personalised and contextualised content to viewers across multiple display deployments and locations is also seen as one way to overcome "display blindness" [Mül+09] – a common issue in public display deployments where viewers stop paying attention to displays in over-saturated environments [Dav+14].

In summary, we are moving to a world with a rapidly increasing number of displays, larger screens and complex eco-systems that increase the number of stakeholders and the range of available content.

1.2 The Need for New Forms of Signage Analytics

The growing number of public displays and digital signs along with the emerging changes in the characteristics of display deployments (e.g. growth in individual display sizes and open networks accessed by a different stakeholders) lead to new requirements for digital signage analytics. It will become essential to understand how viewers interact and engage across a number of digital signs and locations, and measure the influence of digital signs on the viewers' behaviour. Moreover, understanding the effectiveness of public displays and novel applications becomes important to all stakeholders involved including developers, content creators and display owners. We can draw on examples from the history of Web analytics, an application domain that we consider highly related to digital signage analytics due to its focus on user interaction and tracking (in some cases across sites and domains). The insights gained by analytics have revolutionised the Web and enabled administrators to clearly identify potential for improvement of the user experience. The analytics space in the Web is well defined: in 2007, the Web Analytics Association issued the latest version of the Web Analytics Definition as a specification for relevant metrics and aggregations [BBW07]. The standard consists of 26 metrics, of which seven feature the fundamental definition of 'page visits' and 'visitor counts' used as part of the computation of 13 metrics that describe the type of the visit (e.g. referrers, visit duration, landing and exit pages), four characterising the content (e.g. bounce rates or single page visits), and two event and conversion measures to capture user interaction within a single website [BBW07]. Historically, the collection of data relevant for the computation of the described Web analytics reports was performed server-side by parsing and analysing server access logs. In order to enable the capture of single page user interactions and events (e.g. scrolling behaviour, button clicks and cross-device analytics), Web analytics have moved toward the collecting of data in the client device. For example, Google Analytics exclusively features client-side data and user tracking captured through the user's Web browser, and provides developers with the ability to track users' on-page behaviour including customised events such as mouse clicks and scrolling behaviour [Goo18c]. This is an example of the shift of capturing analytics data on the client in contrast to the access logs on the server side, i.e. a form of *user-centric analytics*.

In the digital signage domain the use of novel forms of analytics has not been well studied. New forms of analytics frameworks will become necessary to support large-scale, open networks of digital signage and to allow stakeholders to better understand the effectiveness and success of their displays and applications. Currently, state-of-the-art signage analytics systems focus on the capture and analysis of system-relevant performance measures (e.g. system failures on the signage player [Esp17]), and simple audience analytics such as the number of people walking by a display (e.g. Fraunhofer IIS [Fraa]). Some signage analytics systems provide administrators and content creators with reports on viewer demographics including age groups and gender [Int18]. All of these insights, however, are captured for isolated digital signage displays or in some cases aggregated across a closed commercial signage network [See]. The access to the captured insights in the form of analytics reports is typically limited to a small set of stakeholders such as the commercial display network owners [OnS17].

We believe that it will become crucial for future digital signage analytics to perform a similar shift as that occur in the Web domain from display-oriented analytics to a viewer-centric analytics approach. To emphasise the need for viewer-centric analytics and illustrate the opportunities that will emerge from such a shift in the context of open display networks, consider the following scenario.

> A large number of digital signage displays are located in the entrance area of a railway station that is one of the main transportation hubs for the city. The displays are used to show a mix of digital departure times, advertisements and way-finding content for passers-by. Andreas is a commuter and enters the railway station and proceeds to the platform from which his commuter train will depart. Andreas happens to walk by a public display just at the very time it is showing an advertisement for a specific soda drink. While Andreas is passing by the display, he notices the advert and glances at it due to the prominent location of the display in the entrance area. Even though Andreas was initially heading toward the platform, because of seeing the advertisement Andreas decides to walk to the closest shop instead and purchase the advertised drink. Upon purchasing the drink, Andreas leaves the shop and continues to the platform to eventually board the commuter train. On his way to the platform, Andreas walks by a number of subsequent displays that start showing relevant advertisement and content based on Andreas' recent visit to the shop and his purchasing activity.

We envision that future signage analytics systems will be able to answer the question of potential *cause and effect* from the point at which Andreas entered the railway station and noticed the advert, to the point at which Andreas purchased the product, and the subsequent utilisation of analytics insights to improve Andreas' experience throughout the digital signage network. In order to understand these potential causalities, however, comprehensive analytics are required to be captured, evaluated and subsequently fed back into the sign – providing the opportunity to reliably measure the success and effectiveness of content on public displays and for improving the viewer experience. We divided this problem space into three areas of challenge: 1. analytics data capture, 2. reporting, and 3. automated use of analytics data. To highlight the challenges and opportunities in each of the three areas, we provide a set of challenges for each of the three areas by drawing on the scenario above.

Analytics Data Capture Specific to the scenario, a number of distinct events need to be captured: the content displayed on the digital sign (i.e. which advert was playing when Andreas glanced at the display), the navigation and movement traces of people to understand who is present in the space and how viewers have navigated across displays and locations, glances of passers-by at displays (i.e. Andreas' glance at the display), and their purchasing activities (i.e. Andreas purchasing the advertised product). Such events are typically captured by distinct stakeholders of an open display network (display owners, content providers, retail owners, and Andreas himself) and by heterogeneous systems. For example, Andreas' movements could be captured by an indoor location system, glances at the display through video analytics mounted adjacent to the display, and his purchasing activities through an appropriate shop analytics system.

Reporting In order to understand the potential causalities of viewers (i.e. Andreas viewing the content on the display, and the subsequent visit and purchase activity in a shop), captured events need to be linked together and associated to the individual. However, in the context of open display networks each of these example events are ‘owned’ by a distinct set of stakeholders and captured through systems that are likely disconnected. In order to create insightful analytics reports that highlight the potential causalities, such analytics events form need to be shared across stakeholders, linked and associated with an individual. Additionally, the created analytics reports are equally relevant to all involved stakeholders revealing the effectiveness and impact of digital signs on individuals.

Automated Use of Analytics Data We consider the use of analytics crucial for improving viewer experience with digital signs and the displayed content. It imposes opportunities to use analytics insights and dynamically and automatically adjust the content on digital signs to reflect the insights captured by, for example, providing relevant content to the individual based on a shop purchase or the individual’s preferences. In the context of open display networks, however, a number of requirements (from all stakeholders involved) may compete over screen real estate – making it necessary to develop appropriate systems and interfaces that support resolving such competing requests.

In all cases, overarching analytics are required to understand that the same person caused these events, and appropriate systems and interfaces to support the feed back into the digital sign.

1.3 Research Context

The work described in this thesis has been carried out in the context of e-Campus, the world’s largest digital signage research test-bed established in 2004 and located at the main campus of Lancaster University in North West England [FDE12]. The e-Campus displays deployment consists of over 70 displays and a user base composed of 13,115 students, 4,515 members of staff and a number of visitors. Displays are situated at key locations across the university campus including outdoor information displays along the main university pathway (e.g. Figure 1.2a) student accommodation colleges, departmental and office buildings (e.g. Figure 1.2b), student learning areas such as the library and university-provided learning zones.

The e-Campus displays typically show a mix of static content (including university-wide and departmental news), videos, and websites specifically developed to support public display content. Users can manage displays and content through two Web-based systems: e-Channels [FDE12] and the Mercury App Store [Cli+14]. Both systems were designed to serve the needs of both display owners and content providers. E-Channels can be used to “create content channels – logical containers – for sharing on each other’s displays” [FDE12]. Content providers can place content items (e.g. images, videos and references to Web sites) into channels and share channels with other users of the system. In order to show content from channels on displays, display owners can subscribe their display to one or multiple

(a) Outdoor display located at the university's bus station showing real-time travel information.

(b) Indoor display located inside a departmental building showing a mix of university-related news.

Figure 1.2: Typical display deployments at Lancaster University.

channels that are populated by content providers. For example, displays located in student study zone would be subscribed to multiple channels including university-wide press news and student-facing channels. While in some cases display owners and content providers are represented by the same entity, certain users of the system only create content such as the University press office – similar to application stores for mobile devices. The Mercury App Store introduced the application store model in which content (both static content, videos and Web sites) is modelled as applications to the public display domain. Content providers and application developers can create new applications and offer these through the application store, while display owners can browse the application store and purchase content and dynamic applications for their displays. Similar to e-Channels, display owners can further use Mercury to manage their displays, purchase applications and add these to their displays.

The display nodes at Lancaster run Yarely, an open-source digital signage player developed as part of multiple research projects at Lancaster [Cli+13; MCD15]. Yarely retrieves display schedules from both e-Channels and Mercury through the XML-based "Content Descriptor Set" (CDS) [Cli+13], a file consisting of content descriptions, URLs to the actual static or dynamic content, and scheduling constraints such as date and time availabilities of channels and individual content items. In order to determine which content item to show on the display at any given point of time, Yarely uses a lottery-based scheduling component that considers the options and restrictions specified within the CDS [MCD15]. Both the signage player and back-end components are instrumented with Pheme to capture user interactions and system performance measures – helping content creators, display owners and administrators understand how the signage network performs, which content has been showing and how users have interacted with the displays on campus [Mik+15].

Supporting novel forms of interactions is a unique characteristic of the e-Campus test-bed. In 2017, the entire test-bed was extended to support personalised content through Tacita, a privacy-aware display personalisation architecture [Dav+14]. Over 45 displays have been equipped with Bluetooth Low Energy beacons to detect viewer proximity and support walk-by personalisation (Figure 1.3). Users can download an application for their

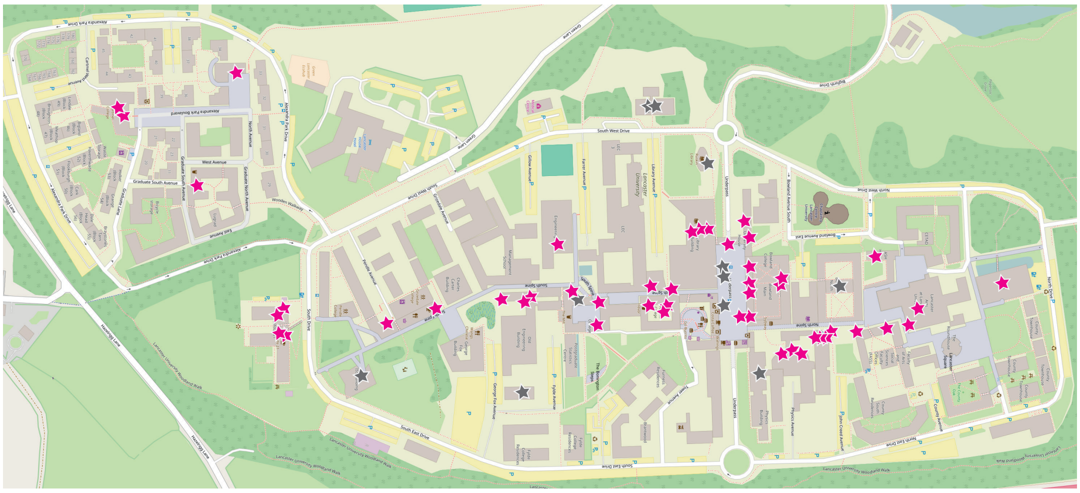

Figure 1.3: Map and overview of e-Campus displays at Lancaster University (grey: standard e-Campus displays; red: personalisation-enabled e-Campus displays).

iOS- and Android-based smartphone devices through which they can express their preferred content and preferences. A variety of back-end systems have been developed and deployed to capture viewer display proximity and send dynamic messages to displays to change the content displayed to a preferred piece of content if a user is in proximity to a display.

A number of components of the e-Campus test-bed have been extended and developed in the context of this work, including the lottery-based scheduling component [MCD15], Pheme analytics framework [Mik+15], and the new Tacita architecture to support display personalisation [Mik+18d]. These components are therefore described in more detail throughout this thesis.

1.4 Contributions

In this thesis, we aim to provide early insights into achieving next generation digital signage analytics whilst preserving the viewers' privacy. We present an exploration into each of the three areas of challenge (*analytics data collection*, *reporting* and *automated use of analytics data*) providing a series of frameworks, components and techniques addressing each of the three areas.

In particular, this thesis makes the following contributions:

C1: Data collection. New techniques for the collection of data relating to individuals' interactions with networks of displays (i.e. across multiple displays and devices), including privacy-preserving approaches. In particular:

1. a framework for categorising analytics data and identifying potential opportunities for measuring the effectiveness of pervasive display deployments and the creation of novel viewer-centric analytics by combining multiple data sources from different stakeholders,

2. an analytics backend system and a set of client libraries to enable us to collect and process display- and application-specific analytics events, e.g. analytics that originate from public display content players and interactive display applications.
3. new insights into the creation of viewer-centric analytics by capturing viewer mobility traces through techniques deployed on the client side,
4. the design and development of a system that supports the capture and processing of viewer mobility data captured on the infrastructure side, and,
5. a novel approach to creating user-centric analytics data and providing entirely new forms of analytics insights in a privacy-preserving way by combining real-world analytics data with synthetic traces of viewer mobility.

C2: Reporting. Novel forms of viewer-centric reports and insights into the opportunities that emerge from leveraging existing Web analytics engines. In particular:

1. the identification and creation of novel viewer-centric analytics reports founded on the combination of traditional display-oriented analytics data and viewer mobility traces captured through viewer-side tracking (C1.3) and generated using mobility models (C1.4),
2. a new set of example reports for novel signage networks that support the delivery of personalised content to viewers, and
3. the identification of similarities between the Web and digital signage analytics and the development and implementation of a mapping of events that span across both domains – allowing us to leverage existing Web analytics engines for the creation of display-oriented analytics reports.

C3: Automated Use of Analytics Data. Novel systems that support the automated use of analytics data on the digital sign to drive content scheduling decisions. In particular:

1. highlighting the need for novel content scheduling systems for digital signs that are able to respond to potentially conflicting content scheduling constraints and requirements in the context of analytics-driven open pervasive display networks, and,
2. the design and development of the first lottery scheduling system for digital signs designed to address the challenge of resolving potentially conflicting scheduling constraints and requirements and, additionally, be able to respond to dynamic content contextual changes from analytics systems.

C4: Evaluation and Trials. We conducted a number of large-scale and long running trials showing the technical feasibility of systems that have been built and deployed in the context of this thesis. Concretely, we showed the feasibility of:

1. the Pheme analytics backend and its integration into the e-Campus display test-bed,

2. the viewer-based collection of display sightings of viewers over a period of six months as part of an in-the-wild deployment of Tacita in the context of Lancaster University,
3. the infrastructure-based collection of display sightings of viewers in the context of a large convention centre equipped with Wi-Fi location tracking capabilities, and
4. the lottery scheduling approach as part of its integration into over 60 e-Campus displays, covering a duration of over two years and 80 million content schedules.

Findings and outcomes of this dissertation have been published in the following peer-reviewed journals, magazines and academic conferences (in reversed chronological order):

1. Mateusz Mikusz, Peter Shaw, Nigel Davies, Sarah Clinch, Ludwig Trotter, Ivan Elhart, Marc Langheinrich, and Adrian Friday. "Experiences of Mobile Personalisation of Pervasive Displays". In: *ACM Transactions on Computer-Human Interaction – TOCHI (in preparation)* (2018)

2. Mateusz Mikusz, Kenny Tsu Wei Choo, Rajesh Krishna Balan, Nigel Davies, and Youngki Lee. "New Challenges in Saturated Displays Environments". In: *IEEE Pervasive Computing* (2018)

3. Mateusz Mikusz, Steven Houben, Nigel Davies, Klaus Moessner, and Marc Langheinrich. "Raising awareness of IoT sensor deployments". In: *Living in the Internet of Things: Cybersecurity of the IoT - 2018*. Mar. 2018, pp. 1–8. DOI: 10.1049/cp.2018.0009. URL: https://ieeexplore.ieee.org/document/8379696

4. Mateusz Mikusz, Sarah Clinch, Peter Shaw, Nigel Davies, and Petteri Nurmi. "Using Pervasive Displays to Aid Student Recall -Reflections on a Campus-Wide Trial". In: *Proceedings of the 7th ACM International Symposium on Pervasive Displays*. PerDis '18. Munich, Germany: ACM, 2018, 6:1–6:8. ISBN: 978-1-4503-5765-4. DOI: 10.1145/3205873.3205882. URL: http://doi.acm.org/10.1145/3205873.3205882

5. Mateusz Mikusz, Sarah Clinch, and Nigel Davies. "Design Considerations for Multi-stakeholder Display Analytics". In: *Proceedings of the 6th ACM International Symposium on Pervasive Displays*. PerDis '17. Lugano, Switzerland: ACM, 2017, 18:1–18:10. ISBN: 978-1-4503-5045-7. DOI: 10.1145/3078810.3078830. URL: http://doi.acm.org/10.1145/3078810.3078830

6. Ivan Elhart, Mateusz Mikusz, Cristian Gomez Mora, Marc Langheinrich, and Nigel Davies. "Audience Monitor: An Open Source Tool for Tracking Audience Mobility in Front of Pervasive Displays". In: *Proceedings of the 6th ACM International Symposium on Pervasive Displays*. PerDis '17. Lugano, Switzerland: ACM, 2017, 10:1–10:8. ISBN: 978-1-4503-5045-7. DOI: 10.1145/3078810.3078823. URL: http://doi.acm.org/10.1145/3078810.3078823

7. Mateusz Mikusz, Anastasios Noulas, Nigel Davies, Sarah Clinch, and Adrian Friday. "Next Generation Physical Analytics for Digital Signage". In: *Proceedings of the 3rd International on Workshop on Physical Analytics*. WPA '16. Singapore, Singapore: ACM, 2016, pp. 19–24. ISBN: 978-1-4503-4328-2. DOI: 10.1145/2935651.2935658. URL: http://doi.acm.org/10.1145/2935651.2935658

8. Mateusz Mikusz, Sarah Clinch, Rachel Jones, Michael Harding, Christopher Winstanley, and Nigel Davies. "Repurposing Web Analytics to Support the IoT". in: *Computer* 48.9 (Sept. 2015), pp. 42–49. ISSN: 0018-9162. DOI: 10.1109/MC.2015.260. URL: http://doi.org/10.1109/MC.2015.260

9. Mateusz Mikusz, Sarah Clinch, and Nigel Davies. "Are You Feeling Lucky?: Lottery-based Scheduling for Public Displays". In: *Proceedings of the 4th International Symposium on Pervasive Displays*. PerDis '15. Saarbruecken, Germany: ACM, 2015, pp. 123–129. ISBN: 978-1-4503-3608-6. DOI: 10.1145/2757710.2757721. URL: http://doi.acm.org/10.1145/2757710.2757721

10. Sarah Clinch, Mateusz Mikusz, Miriam Greis, Nigel Davies, and Adrian Friday. "Mercury: An Application Store for Open Display Networks". In: *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. UbiComp '14. Seattle, Washington: ACM, 2014, pp. 511–522. ISBN: 978-1-4503-2968-2. DOI: 10.1145/2632048.2636080. URL: http://doi.acm.org/10.1145/2632048.2636080

1.5 Scope and Limitations

The work described aims to conduct an *initial exploration* of next generation digital signage analytics in the context of the vision of *open pervasive display networks*. The vision describes future display networks that consist of an increased number of stakeholders contributing both content and displays (*openness*), displays becoming embedded into public and semi-public environments and omnipresent to viewers (*pervasiveness*), and displays becoming interconnected across distinct locations and deployment sites (*networked*). The work at hand has been scoped particularly around the three areas of challenge regarding *analytics data capture*, *reporting* and *automated use of analytics data* (as described in Section 1.4). For each of the areas, we provide a series of data points in the form of frameworks, components and techniques – allowing us to gain a breadth of insights into each area. Additionally, we apply a mixed methods approach that consists of an extensive literature review, and quantitative and qualitative analyses of long-term and in-the-wild trials.

A number of limitations exist to the work presented in this thesis including the following.

- Our work does not aim to provide comprehensive insights in each of the three areas of challenge (i.e. data capture, reporting and automated use of analytics data) but focusses instead on providing a set of *data points* in each of these areas enabling us to capture and report a breadth of insights.

- We did not aim to create a single joined-up end-to-end analytics solution for public display networks but instead provide a set of individual insights, systems and frameworks that can be combined into a common public display analytics system.
- We did not aim to create novel image processing technologies for the capture of analytics data in the context of digital signs (e.g. for capturing gestures) but draw on existing technologies such as Bluetooth Low Energy beacons for location tracking in order to form a foundational dataset for the generation of novel analytics reports.
- The in-the-wild studies and experiments described in this thesis have largely been obtained in the context of the e-Campus display network and with a single deployment in a commercial space. The digital signage testbeds are owned by single stakeholder entities who represent multiple stakeholder groups at the same time (particularly display, space and display owners).

1.6 Thesis Structure

The structure of this thesis is as follows. Chapter 2 (Background) gives an overview of background and previous work in the broader context of analytics, and specifically in the domain of digital signage analytics. Based on the literature, we provide an analysis of the suitability of existing work to support open pervasive display networks. Chapter 3 (Analytics Data Capture and Generation) provides insights into relevant analytics data that can be captured for digital signage and an exploration of the collection and generation of such data sets for the use in new and existing analytics systems. Based on the identified data sets, Chapter 4 (Reporting) outlines concrete examples of the generation and use of novel analytics reports. We demonstrate the opportunities and advantages that result from the shift toward novel viewer-centric analytics. Analytical insights can be used to dynamically inform the kinds of content that is showing on screens and improve the user experience as described in Chapter 5 (Automated Use of Pervasive Display Analytics). Chapter 6 (Trials) describes the evaluation of systems developed in the context of this thesis in both long- and short-term deployments across a number of different trial sites. Finally, Chapter 7 (Analysis, Conclusions and Future Work) presents the contributions and findings of this thesis and highlights areas of potential future work.

Chapter 2

Background

2.1 Overview

We begin the introduction of background and related work by describing models and frameworks developed for the description of different interaction and engagement phases with public displays (Section 2.2). To understand better the current landscape of systems that already support the capture of digital signage analytics, we introduce in Section 2.3 systems that were specifically developed to collect data around digital signs (e.g. for capturing systems performance measures, content logs and the health status of displays), and that provide some insights into the audience (e.g. their demographics and attention levels, and how they navigate within a space). We further introduce public display applications and content (e.g. interactive applications) that enable the capture of explicit interaction events such as direct interaction with the sign, gaze and interaction through mobile phones. To present captured data and learn new insights from such datasets, work has been conducted into the reporting of analytical insights. Section 2.4 presents such previous work including historical approaches originating in Web analytics to describe user navigation patterns through association rules, statistical reports of systems and content, and various visualisations that have been used in the context of digital signage analytics. While analytical insights are often used by display owners and content providers to gain a better understanding about the ways viewers use displays, additional research has been conducted on automatically instrumenting displays and changing their content based on contextual events that goes beyond just simple touch events, as described in Section 2.5. This includes specifically targeted advertising (Section 2.5.1) and the use of analytics to influence the schedule of displays (Section 2.5.2) based on contextual information and viewer presence. Additionally, Section 2.4.3 presents the use of public displays and digital sign analytics in related areas such as retail and the Web. Finally, Section 2.6 consists of a gap analysis in which we clearly highlight the unique characteristics and requirements for future display analytics frameworks in the context of open display networks.

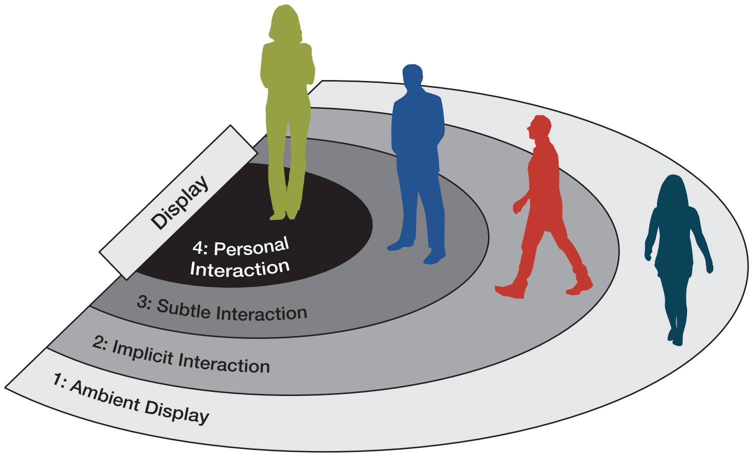

Figure 2.1: Overview of interaction zones by Vogel and Balakrishnan (redrawn from [VB04]).

2.2 Audience Models and Metrics

A number of models and frameworks have been developed for the classification and description of different phases of user interaction and engagement with public displays. Early work in understanding interactions of viewers with public displays was published by Vogel and Balakrishnan in 2004 – with a specific focus on displays that support both interaction from a distance using gestures and direct interaction through a touch sensors [VB04]. The authors developed the *Framework for Interaction Phases* in which they identified four key interaction phases that interactive public display systems should consider when viewers approach displays: (illustrated in Figure 2.1):

Phase 1 – Ambient display: viewers are present in the vicinity of the display and notice the display while it is following its regular content schedule (displays are in their “neutral state” [VB04]).

Phase 2 – Implicit interaction: viewers have been detected by the system to be in the the vicinity of the display. As a consequence, the detection of viewers leads to a change in the content that is displayed without the viewer actively or explicitly interacting with the display.

Phase 3 – Subtle interaction: viewers “approaching the display and providing an implicit cue such as pausing for a moment” [VB04]. Such *subtle interactions* are considered as a basis for a display to reveal personal and contextual information pertaining to the individual viewer.

Phase 4 – Personal interaction: viewers start explicitly interacting with the public display from a short distance through direct touch input or other appropriate input modalities such as gestures.

As Vogel and Balakrishnan point out, viewers “transition from implicit to explicit, public to personal, interaction” [VB04] as they go from phase one (*ambient display*) through to phase four (*personal interaction*) [VB04]. As part of the work, the authors have identified the need to facilitate each of these interaction phases dynamically for an improved viewer experience and provide appropriate content for each phase. For example, the content should change if a viewer dwells in proximity to the screen for a certain amount of time (*subtle interaction*), or if the viewer approaches the display and is close enough to perform direct interaction (*personal interaction*) [VB04]. Vogel and Balakrishnan additionally discuss potential privacy implications of users interacting with displays in public spaces. In particular, the *personal interaction* phase imposes risks of revealing sensitive information. Whilst users can hide some personal information displayed on the screen with their body, Vogel and Balakrishnan note that this technique is only suitable for a certain "class of information" [VB04].

Another example of the exploration of design spaces and interaction zones in digital signage was performed by Rogers and Brignull in 2005 [RB05] – with a specific focus on designing a system that encourages a group of viewers to start interacting with a public display and ultimately to initiate conversations within a group of viewers. The interactive display application that Rogers and Brignull deployed allowed bystanders to share comments on a public display through an input device installed in the vicinity of the screen. The authors observed two interaction patterns: firstly, the application quickly drew in more bystanders and led to a "honey-pot effect" [RB05]. Secondly, people followed a typical interaction pattern in which they first moved into the honey-pot zone, and then started to queue up at the input device to start interacting with the application Rogers and Brignull. The authors further observed that within the "virtual space" that emerged in the immediate vicinity of the display, it "became socially acceptable to spark up conversations with others" [RB05]. The findings lead to the observation that two general interaction zones exist: a highly interactive circle around the display which is visible through the honey-pot effect, and a low interaction zone further away from the display.

Focusing on how passers-by and viewers approach digital signs, Michelis and Müller developed the *Audience Funnel Framework* as a way of describing six typical viewer behaviour and interaction phases as passers-by approach, dwell and leave the immediate vicinity of a digital sign in a public space [MM11]:

Passing by: people who are present in the immediate vicinity and view range of a display.

Viewing and reacting: people who showed *interest* in the display and its content. Such individuals are referred to as 'viewers'.

Subtle interaction: viewers performing "something" that causes a reaction on the display (e.g. approaching a proximity-aware display that detects the viewer and subsequently changes the content).

Direct interaction: viewers who explicitly interact with the display and its content, e.g. through touch or gesture.

Multiple interactions: the same viewers performing *direct interactions* multiple times, i.e. continuously interacting with the display.

Follow up actions: viewers performing additional subsequent actions upon engaging or interacting with the screen, such as following up on the Web [MM11].

While we would expect viewers to go through these phases in chronological order, Michelis and Müller point out this does not necessarily have to be the case [MM11]. Similar to Rogers and Brignull [RB05], Michelis and Müller based their findings on observations conducted as part of a deployment of an interactive application on a public display. In the case of Michelis and Müller, the interactive application mirrored the passers-by virtually on the display as they walked by [MM11]. The framework describes typical ways people interact with public displays – of course, if a display does not offer direct interaction, the phases described might differ. For example, the passing-by and viewing phases would be identical while the direct interactions phase might be interpreted differently in the context of non-interactive public displays.

With regards to the capture and measurement of the effectiveness of advertisements through the means of digital signage, She et al. have proposed a new model for capturing and measuring the "effectiveness of interactive display advertising" [She+14]. The model consists of three main phases in which viewers approach a display:

Attraction: potential viewers becoming aware of a public display in their vicinity, e.g. by spotting its presence or walking by.

Interaction: viewers actively interacting with a display and its content, e.g. through touch or gestures.

Conation: viewers performing *any* actions after interacting with a display.

She et al. defined the "effectiveness" measure of an advertising campaign using public displays as the ratio of the number of viewers of an advert, and the number of people who followed up after the *conation* phase [She+14] which can be compared to the *follow up actions* phase from the Audience Funnel Framework [MM11]. The model developed by She et al. emphasises the importance of considering each of the three phases during the development and deployment of content for public displays – in this case for advertisements.

The knowledge of interaction phases and zones, and the understanding of how viewers approach digital signs are being used to implement novel applications that dynamically react to the presence of viewers depending on the interaction zone a viewer has entered. Ballendat, Marquardt, and Greenberg developed the *Proxemic Interactions* model which, in addition to the proximity of the user to the display, also considers their identity, orientation and movement [BMG10] and builds on top of the interaction transition approach developed by Vogel and Balakrishnan [VB04]. A further publication elaborated on the Proxemic Interactions model and described it as a basis for "the new Ubicomp" [Gre+11]. As part of the model, Greenberg et al. have identified five "dimensions" of relevance for the support of proximity-aware interactions:

Distance: defining either the distance of the viewer to a display or their presence in certain *Interaction Zone*.

Orientation: continuous or discrete orientation of the viewer to the display, e.g. whether the viewer is facing the display.

Movement: both the speed and direction of the viewer while being present in the vicinity of the display.

Identity: unique description of the viewer or entity, i.e. not necessarily the real identity information of a viewer.

Location: relative or absolute location of an entity or the viewer – including relative locations of two entities or viewers to each other.

The authors note that relevant metrics for the support of proximity-aware interactions could be dependent on each other. For example, *Location* as a dimension might be the foundation for capturing data for the remaining four metrics. Equally, each metric in isolation could be enough to feature certain functionalities such as the recognition of an audience [Gre+11]. The Proxemic Interactions model can be used both as a basis for novel application areas for digital signage and for describing viewer behaviour and engagement with an interactive or even non-interactive display.

2.3 Data Capture

The capture of relevant datasets is the foundation to understanding and learning how viewers interact with displays, for example, to determine which content was the most and least effective or interesting to individuals passing-by. In this section, we give an overview of different kinds of datasets and metrics relevant for digital signage analytics – from purely systems-focused performance indicators (Section 2.3.1) to insights into audience behaviour and movements (Section 2.3.3).

2.3.1 Systems

The roots for monitoring the performance and health status of systems, including networks of digital signage and public displays, can be found in the use of the Simple Network Management Protocol (SNMP) [Cas+90]. The protocol was originally developed to support the exchange of "information between one or more management systems and a number of agents" [Sta98] across networks of systems. Since the initial development of SNMP, a number of system and network management tools have arisen [MZH99]. Martin-Flatin, Znaty, and Hubaux have conducted a survey of network management tools allowing developers to chose the best solution for their system [MZH99].

Specifically for monitoring networks of digital signage, a set of commercial monitoring tools have emerged. Esprida is an example for a system that was initially developed as

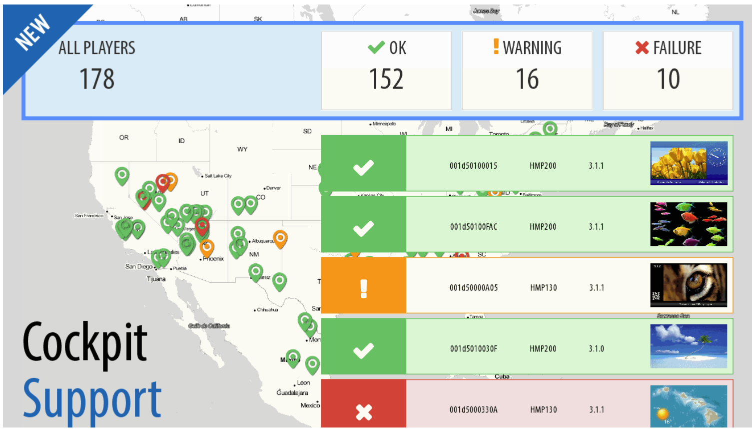

Figure 2.2: Screenshot of SPINETIX Cockpit [Cor].

a traditional remote management tool to monitor and control generic end-point devices – noting that such end-point devices can well represent public displays [Esp17]. While Esprida provides ways to monitor the behaviour of remote devices and digital signage, it is also designed to support the remote management of remote devices. The general motivation of remote monitoring systems is to "offer better service at a lower cost, minimize response time, and maximize the availability and reach of advertising on their networks" [Esp17]. Similarly, OnSign TV is a product that features a real-time dashboard specifically designed to detect "any issues" with the signage network, allowing administrators and technicians to respond fast to potential problems [OnS17].

SPINETIX Cockpit [Cor] is a commercial product that predominantly focuses on enabling display owners to maintain and oversee their remote display deployments [Cor]. As shown in Figure 2.2, the application features the monitoring of remote displays to ensure their hardware status and connectivity in real time. Further, Cockpit tries to automatically detect potential failures and malfunctions, e.g. such as unwontedly disconnected or powered off displays, software malfunctions, environmental problems such as power cuts, or firmware conflicts. SignageLive is a software suite with comparable features to SPINETIX Cockpit and was also designed for the remote monitoring of displays [Rem17]. A Web-based real-time dashboard allows administrators to oversee their entire display network, and view warning and error notifications through the interactive interface. In addition to hardware monitoring, SignageLive also provides an overview of currently and previously played content on displays, and the number of players that have already or are due to update to the latest content schedule.

Traditional signage analytics often uses log files gathered from digital signage as a foundation for the creation of status reports of digital signage players or content played records. An example for such an approach is *CAYIN SuperReporter 2* [CAY] that can be used as an extension to the manufacturer's digital signage players, and collects log files from each individual display remotely to be then used as a basis for the creation of device-centred analytics reports.

2.3.2 Audience Numbers and Demographics

Capturing and counting the size of the audience and their demographics in front of a public display has been a focus both in a commercial context as well as in research. Commercial tools mostly use visual computing techniques to count the number of people in the vicinity of a display and refer to the field as "Anonymous Video Analytics" [Int18; Ham+09; Fraa] (AVA). The name originates from the fact that the sensing and processing of the video feed is performed on or close to the visual sensor instead of in the cloud, and that only the outcomes of the computation are stored instead of the entire video feed. Using this approach, the recognition of the same viewer and thus the computation of unique viewers metrics is not possible [Sla11]. However, "protecting viewer privacy by design" [Cav11] is a clear advantage of using AVA and therefore an important step toward the deployment of such systems that help understand more about the audience without revealing personal-identifiable and other comprehensive insights of individuals. In particular the extraction of demographic information of the audience and retail customers is one of the main selling points of commercial state-of-the-art products. From the use of AVA a set of metrics can be derived: "potential audience" (anyone in vicinity of the display) and "useful audience" or impressions (those who actually glanced at the display) [Sla10b]. Demographic information from commercial signage analytics products that use video analytics include an estimate of the customers or viewers age, gender, and, in some cases, a form of attention and mood measure [Int18; Fraa; See; Qui16; IBM13; NEC13] and ethnicity [GH15]. The overall goal of signage analytics is to help advertisers and content creators quantify the "effectiveness of dollars they spent" [Sla10a] and provide support for targeted advertising on public displays [Far+14; Sla10a].

Intel's *Anonymous Video Analytics* (AVA) [Int18] is a commercial application for processing video analytics feeds in real time with the processing component located on the sensing device. Intel uses a face classification algorithm that identifies people whose heads are facing the camera (and thus also the display), and generates counts and aggregations of the number of visitors of a display. The audience counts are based on the total number of viewers and returning viewers – the application is not able to identify unique viewers. However, AVA does provide insights into the average dwell times of viewers spend to look at the screen. Intel's AVA features an additional set of metrics to the typical audience, impression or dwell time counts to provide transparent insights into the accuracy of the measurements. For example, the *impression count error* describes the accuracy of the current viewer count by comparing the viewer count provided by AVA with ground truth data (e.g. collected through manual counting) – giving insights into the accuracy of the face classification algorithm [San+11]. The system further collects a set of basic demographics on rough age estimates (child, young adult, adult and senior), and gender. Such metrics are included into audience analytics reports as an additional dimension and enable the user of the reports to aggregate viewer and dwell times by demographics.

Researchers from Fraunhofer developed a video analytics system that is comparable with Intel's AVA: Fraunhofer *SHORE* [Frab] is a facial video analytics engine underpinning the

Anonymous Video Analytics for Retail and Digital Signage ("AVARD") product [Fraa]. Similar to Intel, the face recognition is performed on or close to the sensor, and the researchers claim that the algorithm is robust enough to work in differing lighting conditions, even when utilising consumer-grade video cameras. While digital signage is one application domain, the software is advertised to be used in retail, health and other application areas. SHORE [Frab] goes one step further than just retrieving basic demographic information and is additionally capable of determining the emotional state of viewers purely through video analytics techniques. The emotional state is categorised in terms of four facial expressions: *happy*, *sad*, *surprised*, and *angry*. Fraunhofer SHORE additionally provides a fine granular age estimation and returns an actual age instead of an age range – including a deviation metric to communicate the accuracy of the estimated age. This product is the foundation and visual analytics engine for Fraunhofer's *Anonymous Video Analytics for Retail and Digital Signage* [Fraa] software suite, which provide an example use case SHORE and the supported metrics and reports specifically in the digital signage domain.

The detection of emotional states and moods is not a unique feature. Other commercial visual analytics products have been specifically developed for the digital signage domain that use a similar approach for face classification. Quividi [Qui16] provides similar visual analytics tools that are capable of detecting a number of metrics that are in common with Fraunhofer SHORE: "opportunities to see" the content (i.e. counting people who walked by the display but have not glanced at the screen), number of viewers, dwell and attention times, gender and age estimates, attention states and moods "from very unhappy to very happy" [Qui16]. As a unique feature, Quividi additionally supports the detection of facial attributes including facial hair, glasses and sunglasses [Qui16]. Specifically designed for analytics for kiosks, Meridian uses video analytics to detect and count "potential users" and actual users and the retrieval of the collected data in real-time, e.g. for the use of interactive and adaptive display content [Sla16]. The classification, however, is limited to age and gender, though could be extended with other video analytics products and enriched with interaction logs captured directly through the interactive kiosk software. *SCALA Advanced Analytics* [Sca] even allows the plugin of a range of sensors and actuators that can be individually programmed and dynamically change their behaviour based on audience presence and viewer counts. For example, displays could change the content displayed based on an approaching audience or interactions in proximity to the display [Sca].

A broader approach is used by IBM *Intelligent Video Analytics* [Ham+09; IBM13] and NEC's *FieldAnalyst* [NEC13] software. While the previous products focused mainly on face recognition classifiers and required the camera to be mounted on the screen, IBM focuses on analysing video feeds from CCTV cameras [Ham+09]. Similarly, the *FieldAnalyst* software captures faces and people from video streams and is capable of measuring a basic set of demographics (age, gender, distance to the screen, and viewing time), the number of viewers of a display and, additionally, the number of entrances and exits in a space without the need to place the camera at the display [NEC13]. Similar to Fraunhofer AVARD [Fraa], FieldAnalyst is designed for the digital signage and retail domain and provides ways for "target analysis" and “non-buyer” analysis – helping display owners and content providers to understand which user groups are engaging with displays. Seemetrix is able to return a similar set of metrics: it consists of the capability to capture a rough age and gender classification of viewers [See]. Reports are extended by an attention measure per viewer which is calculated from the total duration a viewer has spent looking at (or in the direction of) the display, i.e. the duration in which the viewer has been “attentive” [See].

While commercial products focus on providing an end-to-end system for capturing and reporting information about the audience, one of the main focuses in the development and use of visual analytics tools is to answer questions about the user behaviour and the ability to track individuals [LCK13]. Of course, systems that support audience tracking in the context of pervasive displays are also capable of generating audience numbers.

Examples of specific signage analytics work include the analysis of pedestrian traffic around a public display performed by Williamson and Williamson [WW14]. The authors placed a video camera on the display and used visual computing techniques to both count and track people walking in the surrounding area of the display deployment. While the focus of the system was to track people, the same approach could be used to simply count the number of people who are in the immediate vicinity of the display and produce an audience count measure. Using a depth-camera mounted to the display and facing the audience as a source, Tomitsch et al. deployed a public display to conduct a study to understand the level of care and attention of viewers toward content that is shown on the displays [Tom+14]. The video stream was recorded as part of the deployment and the authors were able to use it for a better understanding about the audience (including the number of people) and their behaviour in front of the display. In a similar approach, Farinella et al. equipped a public display with cameras and developed a system that supports the identification and recognition of returning viewers at a public display based on biometric features [Far+14]. Parra, Klerkx, and Duval used visual computing techniques to automatically generate an audience count of people passing by at an in-the-wild deployment at Brussel’s largest train station [PKD14].

In addition to simple audience counts, visual analytics based systems are often also capable of capturing the user dwell time in proximity of the display, and their view times of the display and content [RS13]. More recently, Elhart et al. published the “Audience Monitor” – a toolkit specifically designed to count the number of people approaching a display and their dwell time [Elh+17]. Utilising a mix of different sensing technologies, Gillian et al. developed “Gestures Everywhere”, a system that is able to track an individual across multiple displays through a number of sensing technologies such as Bluetooth Low Energy beacons and video cameras [Gil+14]. In addition to providing context-aware content to the viewer, the system also supports the tracking of individuals across multiple displays and locations and serves as a basis for the generation of analytical insights such as audience counts.

Whilst the presented work typically requires the use of video cameras mounted at the display facing the audience, other video analytics tools utilise surveillance cameras that capture a broader view of the vicinity of the display. IBM Intelligent Video Analytics uses such an approach in which it is possible to search for specific faces, the extraction and filter for detail demographics, however, is not possible [IBM13]. Note that, for the purpose of conducting and analysing research experiments, demographic information of an audience have often been conducted manually through observations, e.g. in [Alt+11a] and [PTK18].

The use of video analytics techniques in order to capture audience numbers and demographics can potentially impose a privacy risk to individuals present in front of the display. Previous work, however, has developed approaches that address the potential privacy risks. For example, Intel AVA [Cav11] conduct the analysis of the video feed close to the sensor and report the generated numbers (e.g. the number of viewers engaging with the display) instead of the video feed. Similar approaches are taken with the concept of "Edge Analytics" in which computations are performed on the edge of the cloud close to the sensor both for performance reasons and for privacy preservation [Sat+15].

2.3.3 Audience Engagement and Movement

2.3.3.1 Proximity-Aware Systems

The use of information on the proximity of viewers to displays has been explored in research mainly in the context of the development of novel interactive applications that change their content or enable certain interaction modalities based on approaching viewers or the distance of viewers to a display [MG15].

Vogel and Balakrishnan have developed a prototype system to specifically support the transition "from implicit to explicit, public to personal, interaction" [VB04] of viewers with a pervasive display. This work is an early example of a proximity-aware pervasive display system that automatically adapts the content and interaction modalities to the viewer based on their distance to the screen. To measure the engagement and proximity of viewers, the authors chose active markers and motion detection sensors (the state-of-the-art at the time at which the paper was published). The distance measure of the viewer to the display was used to dynamically adapt the content and enable certain interaction modalities. For example, as soon as a viewer has reached the immediate vicinity of the display and is close enough to perform direct interaction, the display enables the touch input sensor and changes the content to an interface that is suitable for direct interaction [VB04].

More recently, Wang, Boring, and Greenberg developed the Proxemic Peddler – an interactive application as a demonstrator for an "advertising display that captures and preserves the attention of a passers-by" [WBG12] by utilising proximity information and tracking individuals while they are in the vicinity of the display. The authors describe one of the key advantages in monitoring the immediate vicinity of a display as the ability to tailor the content to the viewer and their current location [WBG12]. The demonstrator continuously captures the movements of individuals and their "attentional state" in real-time and adjusts the content accordingly, e.g. by showing "rapid animations" to viewers passing by, and more detailed content to viewers that decided to dwell in front of the display [WBG12]. Wang, Boring, and Greenberg work builds on top of the Proxemic Interactions model developed by Ballendat, Marquardt, and Greenberg [BMG10] (described in more detail in Section 2.2) and the Proximity Toolkit developed by Marquardt et al. [Mar+11]. The tool-kit is a software that incorporates the interaction model and enables developers to create interactive applications that utilise proximity information. The tool-kit supports the collection of relevant analytics data in accordance with the Proximity Interactions model: *distance*, *orientation*, *movement*, *identity* and *location*. In addition to sensing data about individuals, measurements about the (spatial) relationship between two individual entities (e.g. their distance, angle and velocity difference) and the pointing relationship between two points (e.g. intersection) is collected and made accessible by the tool-kit API [Mar+11]. A number of analytical insights can be derived from such data. For example, the speed and angle in which viewers are approaching the display, and their potential relationship to other bystanders can be captured and used to measure changes in behavioural or movement patterns depending on the displayed content.

The use of proximity data has also been explored for the design of interactive and collaborative spaces. Dostal et al., for example, have designed a system to support "collaborative proxemics" [Dos+14]. The authors utilised a combination of Microsoft Kinect sensors and consumer RGB depth-cameras to support, among other features, the tracking of multiple users – with a limitation to up to four users in real-time. In addition to proximity information, the tool-kit also gathers data on the proximity to other (potential) group members, and the attention level of the users of the system. The system is purely based on computer vision, and the authors included a confidence level measure that transparently indicates the accuracy of the results. Gillian et al. use "Gesture Everywhere", a system that captures gestures and audience presence information through depth-cameras in real time at MIT Media Lab to recognise the proximity to the display of individual viewers and select an appropriate piece of content [Gil+14]. While the data is initially captured within the analytics module, the system makes the data available for applications and content running on the display and allowing the immediate use of current gestures and audience engagement for changing the content, and viewers to interact with the display. An opposite approach was taken by Alt et al. with the GravitySpot deployment [Alt+15]. Instead of tailoring the content to the viewers location, the authors have instead tried an approach in which the content on the display is used to guide the viewer to the "sweet spot", i.e. the best location in which the user can consume the content displayed [Alt+15]. Visual analytics techniques are used to track the current location of the viewer and increase or reduce the blurring on the display until the viewer has found the desired location [Alt+15].

The utilisation of proximity-aware interaction models is not restricted to the research environment – a number of commercial providers offer pervasive display systems that are based on user proximity information [MG15]. Novo Ad, for example, is a product that is typically deployed in public bathrooms. Sensors are used to detect a person standing in front of a mirror to then automatically display advertisements on the mirror or in front of urinals without allowing the user to opt-out or turn the system off [Cap17; MG15]. Targeting touch-enabled kiosks, Meridian is a commercial product that uses visual computing techniques to detect an approaching audience and then automatically change the content to provide incentives for the user to start explicitly interacting with the display [Sla16]. Proximity data is further collected to provide analytical insights about the performance of the kiosk to the display owners. Measuring the distance of a viewer to the display is a common metric in digital signage analytics. Quividi, for example, captures such data for each person in vicinity to the display [Qui16]. Such data is made available in real-time and Quividi explicitly mentions the use for proximity-aware systems and the adaptation of content based on audience presence and proximity as one potential use case [Qui]. Similarly, Cayin provides an interface for the plugin of external sensors (such as proximity sensors) to enable “interactive and personalised applications” and “multi-functional services” [CAY].

Greenberg et al. identified a set of “dark patterns in proxemic interactions” [Gre+14] in the context of interactive and context-aware public displays. In particular, the authors highlight issues with regards to user privacy – for example, displays revealing a viewers interests by providing personalised content as they approach the display [Gre+14]. The provision of personal content through public display has been further considered by early work such as Shoemaker and Inkpen [SI01] who developed the *Single Display Privacyware* system designed to provide personalised content in the context of a shared public display.

2.3.3.2 Capture of In-Door Viewer Movement Patterns

In order to measure the influence and therefore also the effectiveness of digital signage and pervasive displays we need to be able to track the audience beyond and across single displays within closed environments. This includes the need for more detailed understanding of viewer navigation patterns, i.e. focused on the viewer instead of the physical device of a display. In related domains such as activity recognition and human sensing such shifts were previously described as “human-centric personal analytics” [LB14] and “human-centric sensing” [SAS11].

Tracking and the required localisation of individuals indoors is a challenging task but equally is an “important source of context for ubiquitous computing systems” [WH08]. An early example for determining the indoor location of users dates back to 2004 in form of a patent by Gray, Jeffrey, and Chery [GJC04]. The general idea is to use triangulation in spaces densely equipped with wireless access points, and identifying individuals by their mobile device’s MAC addresses. While the computed location is often an approximation of the actual location of the person or device, the movement traces acquired over time are still informative enough to give insights into certain navigation patterns of an individual and are certainly accurate enough to, for example, determine in a shopping mall scenario which shop a person has entered and how long they have stayed inside [Nan+13]. Researchers have also developed improvements to WiFi location tracking. This includes the work of Woodman and Harle who focused on the development of an indoor localisation system that, in combination with Wi-Fi location traces, uses movement data captured by foot-mounted units [WH08].

The use of Wi-Fi- and infrastructure-based location tracking has also been explored and deployed in commercial spaces. The “retail store of the future” [Air13] describes the use of customer movement tracking to improve staffing, and also to provide better marketing and even perform A/B testing inside the store [Air13]. Generally, the combination of multiple data sources have a high potential and will be discussed in the subsequent sections. A similar approach has been used to develop one of the largest and longer-running test-beds for indoor localisation – which has been deployed at three different environments: a university campus, resort island and a convention centre [Kha+13; Jay+16]. The location data has been made available as a platform and service to other researchers and developers, enabling the utilisation of the data for the development of novel applications and data analytics [Jay+16]. The use of Wi-Fi location tracking, however, brings a number of tradeoffs and considerations. For example, the performance of location tracking algorithms, and therefore also their accuracy, may be highly dependent on the occupancy level in the space that may be changing significantly within small time frames [MB13]. Khan et al. conclude that Wi-Fi- and infrastructure-based location tracking for indoor environments still remains a challenging approach [Kha+13].

The use of mobile devices and smartphones are key for enabling indoor location tracking. While Wi-Fi fingerprinting is one approach, built-in smartphone sensors offer a wide range of opportunities for collecting information about the context in which the user is present [BH13] – including for determining their spatial location. Park, Shin, and Cha have used the magnetometer and accelerometer of smartphones to automatically detect each time the user turns around a corner or changes their direction in real-time [PSC13]. The authors mapped detected corners and turns back to the spatial layout of the building and were able to compute the trajectory of the users movement path. Park, Shin, and Cha note, however, that this approach requires a certain layout of the space that needs to be manually configured within the system and might only work with buildings of a certain layout (e.g. with clear corners) [PSC13]. In a related approach, Roy, Wang, and Choudhury designed an application that uses smartphone sensors to determine the person's walk direction [RWC14]. A combination of compass, accelerometer and other sensors are used to compute the direction in which a person is walking – such information can be then mapped back onto a spatial map and therefore used to draw conclusions on how viewers navigated throughout the space. Indoor localisation tools can benefit from such data to, for example, improve the accuracy of the location determination [RWC14]. Even light sensors can be used to accurately determine the indoor-location of a person in a controlled environment as shown by Hu et al. [Hu+13]. Researchers have also explored approaches in which the use of active or passive tracking devices is not required. One example of such systems is WiTrack which measures the body reflections of radio signals [Adi+14]. Adib et al. point out the relatively high accuracy which, at the time of publication, even exceeded the accuracy of other radio frequency tracking systems that required the use of transmitters. To support tracking through radio signals, however, the environment needs to be heavily equipped with appropriate tracking devices. In its current version, the system can only track a single person at a time present in the space.

Capturing (indoor-) location traces can potentially reveal sensible insights about individuals and could potentially be misused to monitor individuals, for example, in the context of a work environment [Wan+92]. Early work such as the Active Badge system designed to track individuals indoors already considered privacy as an important issue [Wan+92]. The authors note that in the system designed devices rather than individuals are tracked, allowing individuals to interrupt the location tracking by putting away the device at any time – i.e. providing full control over the systems' insights. Want et al. further note that location tracking "technology is rarely inherently bad; it is just that it can be used for good or bad purposes" [Wan+92].

2.3.3.3 Large-Scaled Viewer Navigation and Movement Patterns

Moving out from closed and indoor environments, an understanding of how viewers navigate beyond the immediate vicinity of a digital sign is essential to answer questions on the display effectiveness and potential behaviour and movement changes due to interactions with the sign. Linking back to the motivating scenario (see Section 1.2), this level of tracking is key to link the viewer path from the display to the nearest shop – and therefore answer the question about a potential change in behaviour due to seeing the display content. Visual analytics and image processing have been widely explored in research as a common approach to track people behaviour. Candamo et al. have brought together an extensive literature survey on algorithms and tools that use image processing to extract the behaviour of individuals or groups from a video stream [Can+10]. Example applications of such systems include automated behaviour detection, interactions of people with each other and even automated fraud and safety detection [Can+10]. An early example for the development of a visual analytics software for the activity recognition of customers in a retail environment was performed by Krahnstoever et al. [Kra+05]. The authors developed a novel approach in which RFID technology was used as an additional source for the identification of objects on a shelf which have been picked up by the customer – while video analytics was used to capture the customers' movement traces. The resulting analytics provided novel insights into the ways customers moved across the store and interacted within the retail space, on a level similar to user interactions on an e-commerce website (e.g. the tracking of products that were added and removed to the virtual or physical shopping basket). This could be linked to digital signage the customers may have interacted with before picking up and purchasing a certain product in the physical store.

Generally, video cameras and the subsequent video analytics computation are a common technique to use for the extraction of detailed information about people including their body attributes such as hair type, eye wear and clothing colour [Ham+09]. The authors mention the recognition of the same person across multiple video feeds as an explicit use case of this system – thus allowing the collection of movement patterns in a large scale and across locations as an alternative to using face recognition systems. The tracking of individuals across multiple video sequences and locations has also been explored by Yang et al. [Yan+07]. While their work is focused on the development of the actual visual computing algorithms that are used to track an individual across multiple frames and angles, the resulting insights allow the tracking of an individual across locations without the need for an active or passive tracking device. Of course, while the primary use of such a system is often motivated by safety arguments for the automated detection of fraud or crimes, it can also be used for capturing analytical insights relevant for the digital signage domain. The computation of behaviour and navigation patterns of individuals after passing a digital sign could allow administrators and providers to gain insights into the effectiveness and performance of a display or content displayed. The recognition of people across multiple video sequences is a difficult problem. Mitzel et al. designed a system that uses sparse detection and segmentation to follow an individual, and supposedly is robust enough to handle occlusions and multiple camera feeds [Mit+10]. The approach of using additional data sources has also been used for improving the accuracy of people tracking. Teixeira, Jung, and Savvides combine visual analytics from existing surveillance infrastructures with on-device sensors (in this case, accelerometer and magnetometer) to uniquely identify people through their mobile phone's unique identifier [TJS10]. This allows the recognition of the same person even after returning to the monitored area. Of course, this approach requires the person to carry a smartphone with a dedicated application as an active tracking device. As an example for the use of behavioural analysis, Girgensohn, Shipman, and Wilcox have developed a system that uses visual analytics to determine and track activity patterns in a retail store [GSW08]. The tool produces heat maps of common movement patterns and frequently visited places, and the visualisation also includes the location and speed of people. This can help in understanding the kinds of activities customers were performing in the space [GSW08].

Using visual analytics for activity and behaviour recognition and tracking is a common use case in research as identified in a survey conducted by Candamo et al. [Can+10]. Alternative approaches for tracking people to gain more insights into their navigation patterns, for example, in the context of adventure and amusement parks, involve the use of active GPS tracking devices [RCS10]. Konidala et al. use mobile applications to collect additional information about visitors, including their location coordinates as they move around the park [Kon+13]. The authors point out the usefulness of such data for analytics purposes for park providers. On an even larger scale, cellular data can provide movement traces of individuals [Bec+13] – though the granularity does not appear to be suitable for the use in the context of signage analytics.

The detection of people and objects in video sequences, however, is a complex problem. Researchers have worked on methods to improve the accuracy and performance of such systems, e.g. by combining computer vision with “models of pedestrian dynamics” [Ant+06]. In addition to the technical challenges in audience tracking through video analytics, other challenges and concerns arise. The use of video analytics and tracking of individuals has been pointed out to be privacy invasive and a number of concerns have been raised within the research community (e.g. context-aware displays that recognise individual viewers make their private information visible on public displays) [Gre+14]. To address such issues, researchers have worked on methods to enable the use and collection of such insights while still preserving individuals privacy. Zhang et al. [Zha+10], for example, discuss privacy issues of closed-circuit television cameras and provide a privacy-preserving solution for pedestrian tracking and recognition. Instead of storing an image of the viewers face (i.e. a frontal image of an individuals face), the system computers biometric features of the face and stores a hash of the biometric features [Zha+10] – similar to how systems currently hash password strings to avoid the storage of such sensitive information in plain text. This approach prevents potential attackers from decoding facial and biometric information stored on the device while allowing the system to recognise reoccurring viewers.

Viewer navigation paths (whether captured via video analytics or location tracking systems such as GPS) represent highly sensitive data and may violate the privacy of individuals. A number of systems, however, have been developed that address issues of privacy in this context. LocServ is an early example of a system developed by Myles, Friday, and Davies [MFD03] specifically designed to allow individuals take full control over location data captured and processed by defining and applying policies. For example, users can express times and contexts in which location tracking is acceptable (e.g. while at work) whilst rejecting location tracking in other contexts (e.g. while at home) [MFD03]. The system has been partially motivated by pawS ("a privacy awareness system for ubiquitous computing environments") [Lan02] supporting the implementation of 'data usage policies' allowing users to both express their preferences for the usage of personal data and track the usage of it.

2.3.4 Interaction Events

Understanding more about interactions and engagement with digital signage and content showing on displays is a key aspect in drawing conclusions about the effectiveness and usefulness of both display deployments and content [Alt+12]. Multiple ways exist in which users can interact or visually engage with a public display. She et al. conducted a survey to extract common interaction modalities in the digital signage domain and identified *presence*, *direct touch*, *gesture* and *remote control* as a set of interaction categories [She+14]. For each of these categories, different software- and hardware-based capture techniques are available to capture the data that is relevant for signage analytics. Whilst we described proximity-aware systems that take viewer presence into account in Section 2.3.3.1, we will give an overview of related work in which researchers have developed and deployed digital signage systems that allowed viewers to interact through remote controls in form of mobile phones, gestures, touch, and, as an additional category, gaze. All of the described technologies could function as data capture techniques and be used to build up analytics-relevant datasets about the interaction and engagement of viewers with digital signage.

2.3.4.1 Touch

Direct touch may be a likely choice for the deployment of interactive public displays such as kiosks (e.g. the UBI Hotspots deployment in Oulu, Finland [Oja+12]) and offer a range of opportunities for the collection of analytics-relevant data. As part of the proximity-aware system developed by Vogel and Balakrishnan and initially described in Section 2.3.3.1, the authors included direct interaction for the "personal interaction" phase, i.e. when users have approached the display close enough [VB04]. The authors describe the higher accuracy and the occlusion of the content a user has requested from the display as two among other advantages of direct touch over other interaction modalities such as gesture [VB04]. The use of multi-touch interfaces in which multiple users can engage with the display at the same time infers new challenges. Jacucci et al. have specifically looked at the design of application that support such multi-touch interfaces [Jac+10]. From an analytics perspective, the particular focus lies on the differentiation between multiple users – which has also been identified as one of the design challenges in this work in the context of supporting “parallel interaction” [Jac+10]. For the evaluation of their multi-touch public display prototype, the authors have, in addition to surveys, recorded viewer interactions with a camera located in the vicinity of the display and used the recordings to perform an observation of viewers.

In contrast, Alt et al. have chosen a different approach to analyse user engagement with the deployment of Digifieds, an interactive classified application for public displays [Alt+11a]. The authors made heavy use of log files collected on the server to gain insights into the number of users, usage times and the kinds of classifieds and content users have requested through direct touch interactions on the display and additionally through a mobile phone client. These touch events and content requests were used to identify peak usage times throughout a day and content from which users have requested additional information [Alt+11a]. In addition to log files, Alt et al. have also performed observations and controlled studies to understand how viewers engage with the application.

The use of touch modalities can have potential impacts on viewer privacy particularly around the interaction with display applications. For example, Kim et al. [Kim+10] identified the vulnerability of touch interfaces (in the context of tabletops) with regards to ‘shoulder surfing’ (i.e. bystanders capturing interactions and potentially sensitive information revealed by the viewer such as authentication). Work that addresses issues of privacy in the context of touch interactions on public displays includes Sharp, Scott, and Beresford [SSB06] who developed a system that allows viewers to use mobile phones instead of touch-based interfaces on the display in order to enter sensitive data.

2.3.4.2 Gesture

Depth- and infrared cameras, typically mounted at or behind the digital sign, have been a prominent choice of technology among researchers to enable gesture-based interactions. For example, in 1997, Matsushita and Rekimoto developed the “Holo Wall”, a large display wall composed of a glass wall with a rear-projection sheet, a projector and infrared-camera pointing from behind the wall at the audience to capture their gestures and interactions [MR97]. More recently, Peltonen et al. used a technologically similar setup for capturing viewer gestures and to track the audience in front of the “City Wall” [Pel+07]. Some of the main features of the system include “hand posture and gesture tracking” and “computer vision based tracking” [Pel+07] – both approaches can be used in combination with the infrared-based camera to log user interactions through gestures. With the release of Microsoft Kinect and an appropriate software development kit, depth cameras became a widely used and powerful tool in the research community and are used even beyond the digital signage domain, e.g. for object recognition and human activity recognition [Han+13].

Müller et al. deployed a public display using a Kinect sensor to capture user interactions and engagement [Mül+12]. The authors followed the recommendations of the Audience Funnel (see Section 2.2) for the implementation and deployment of the display. The proximity information about the audience captured by the sensor is used to change the content shown on the display in order to make the passers-by aware of the gesture-enabled display and to encourage them to start interacting. Using sensing technology to make the audience aware of gesture-enabled displays has also been explored by Walter, Bailly, and Müller who conducted experiments with different kinds of gestures, and strategies of revealing gestures to the audience [WBM13]. For this purpose, a similar approach to [Mül+12] was used: a depth-camera was deployed in conjunction with a public display and used to both detect approaching viewers and trigger content when a viewer was detected, and provide gesture interaction support [WBM13]. In general, Walter, Bailly, and Müller have found that gestures appear to be the better technique of interaction for viewers (compared to, for example, touch interfaces). However, in accordance with the Audience Funnel, the authors identified the importance of communicating to the viewer that a display is gesture-enabled as an important step. Interestingly, the authors found that instead of the display dynamically changing its content as someone approaches the display, a permanent information should be displayed to make people aware of the gesture functionality [WBM13].

As an extension to standard digital signage and the use of depth cameras, more sophisticated setups have also been explored for the support of gesture interactions. For example, Sodhi, Benko, and Wilson developed a system that uses a combination of projectors and depth-cameras to project directly onto the person's hand. In this case, the hand acts both as a gesture and display device, and the projection on the hand is used to communicate supported gestures to the user [SBW12]. Data collected from the depth-camera is therefore used to provide instant feedback to the user, and can also be collected for further analytics purposes. In the context of a retail store, Ravnik, Solina, and Zabkar have explored how audience analytics data can be used to predict customer behaviour by utilising machine learning algorithms and video analytics [RSZ14]. The developed system was able to predict in over 70% of cases whether a customer will perform the purchase of a product upon the interaction with a digital sign. The authors note that such an approach "can also be used to predict the roles of an individual in the purchase decision process" [RSZ14] – therefore building a potential foundation for novel insights and interactive displays. Equally, more simplistic techniques can be used for capturing audience attention and supporting gestures on public displays. Hardy, Rukzio, and Davies developed an application for public displays that uses a consumer-grade web-cam in combination with motion recognition software [HRD11]. While gesture recognition was successfully implemented, the authors have based their data analysis purely on observations and manually collected and measured attention levels, glance times, and other information about viewers such as demographics – not using the insights that might have been captured by the system.

Similarly to touch-based interactions, gestures can be subject to privacy concerns by viewers. Clark and Lindqvist [CL15] identified a number of threat models particularly around the use of gestures for user authentication. For example, potential attackers may capture gestures performed by individuals simply by observations (as gestures in the context of public displays are performed in public or semi-public spaces). Certain counter measures, such as the increase of gesture features detected by a system can increase both security and privacy [CL15].

2.3.4.3 Gaze

Gaze tracking has been explored in the context of digital signage in two domains: enabling support for gaze as an interaction modality [Ead+04; SHT10; ZBG13], and for capturing analytics data about prominent regions on a public display [Bur+05; Sip+10; ZBG13].

An early example of the use of gaze to enable interactions on public displays is Eye-Guide [Ead+04]. Motivated by the lacking support of public displays to select and retrieve sensitive and personalised information, Eaddy et al. utilised a head-worn eye-tracking device (in form of a separate head- and eye-trackers) in combination with a headphone and microphone to allow the user to select content on the display without specifically pointing toward it (e.g. the destination on a map to which the user requires guidance) [Ead+04]. The main focus of this work was on the support of novel interaction modalities and the support for privacy-preserving ways of retrieving information from public displays. More recent approaches to support gaze have shifted from the use of user-worn eye-tracking to tracking devices that are mounted on the public display pointing toward the audience. San Agustin, Hansen, and Tall, for example, showed that supporting eye-tracking for viewers can be achieved with off-the-shelf consumer components, in this case a combination of a video camera and infrared light sources [SHT10]. However, such systems can have certain limitations: the user is required to go through a short calibration process prior to being able to interact with the screen, and the distance from viewer to display cannot be greater than 1.5 meters [SHT10]. To provide immediate support for interactions without the need to go through a calibration phase, Zhang, Bulling, and Gellersen developed an eye-tracking system for public displays with simplified interaction capabilities: viewers can use their eyes to indicate three gaze directions: left, right and centre [ZBG13]. In contrast to other systems, Zhang, Bulling, and Gellersen use a single consumer camera without additional infrared light sources and other sensing technology. The developed system is further able to detect "gaze attention to the screen" [ZBG13], i.e. capture whether a viewer present in front of the display is explicitly looking toward the display. The authors emphasise the system to be "calibration-free and person-independent" [ZBG13], making it therefore suitable for public display deployments to capture insights about audience behaviour.

The use of gaze in the context of analytics has inspired work investigating Web advertisements such as the research conducted by Burke et al. who specifically looked at "banner blindness" [Bur+05]. As part of the controlled experiment conducted by the authors, the data collected included the directions in which users have looked, and provided novel insights into the areas of websites that are focussed by users [Bur+05]. The emergence of Google Glass and similar devices enabled Rallapalli et al. to take the eye-tracking approach into the "real world" outside of an controlled lab to capture "physical browsing by users in indoor spaces such as retail stores" [Ral+14]. The collected data included dwell times, gaze directions at products, and reach out for certain items [Ral+14]. The authors achieved this high level of tracking through a combination of sensors on Glass and the user's smartphone. As Rallapalli et al. point out, the resulting dataset can be best compared with the data that could be captured by the user browsing through an e-Commerce website and viewing certain products – in this case, without any data collection by the retail store. Sippl et al. developed a prototype system that was specifically designed to track and collect the viewer's gaze direction in the context of public displays to support analytics [Sip+10]. The authors emphasised the importance of understanding the gaze direction of viewers specifically for advertising enabling dynamic content to focus on the region of a display on which a viewer is the most interested or attentive [Sip+10]. Similar to the system developed by Zhang, Bulling, and Gellersen [ZBG13], Sippl et al. use a single camera mounted on the top of the display and are able to detect general regions on the display that are in focus by the viewer (i.e. top-left, top-right, bottom-left and bottom-right) without the need for calibration. Gaze data are collected in real-time enabling the dynamic integration to other applications beyond targeted advertising [Sip+10].

Gaze-tracking is considered a privacy-preserving and secure modality for viewers to interact with public displays and other devices. De Luca, Weiss, and Drewes [DWD07], for example, developed a system that utilises eye-gaze in order to support the entry of passwords for cash machines. Due to the subtlety of eye-gazes, it is significantly more difficult for bystanders to capture the input and interactions of viewers [DWD07].

2.3.4.4 Mobile Phone

Researchers have also explored the use of mobile phones for interacting with public displays – including the use of mobile phones as pointing devices for direct interaction. "Digital Graffiti" is a system developed by Carter et al. [Car+04] that allows the interaction through a personal mobile device to allow users the sharing of personal drawings on public displays. In this case, gestures are captured through a dedicated application running on the user's mobile device which were send to the display explicitly by the user. A similar system was developed by Scheible, Ojala, and Coulton for *MobiToss* [SOC08] to allow users the capture of short videos. To release the video to a public display nearby, users were required to perform a "throwing gesture" [SOC08]. Similar to [Car+04], gestures were also captured directly on the device and used to trigger the video upload. While the focus of the authors was not on the analytical use of such gesture data but on providing a novel user experience, this work yet shows early opportunities of how data can be captured client-side (i.e. supplied by the viewer) in the context of digital signage in addition to the sensing that can be performed on the sign through visual sensors – similar to the concept of client-based event tracking in modern Web analytics that are sourced from the user's browser [SOC08].

The use of mobile phones has also been explored to support implicit interactions of viewers with public displays. Davies et al. designed a system to support the personalisation of public displays [Dav+14]. Viewers can express their preferences through a dedicated mobile phone application which then monitors user locations and automatically requests content changes when the user is close to a display [Dav+14]. The system is designed to support content changes even if the user is not explicitly using the Tacita mobile application. This functionality is achieved through the location monitoring in the background compared with a map in which the developers have specified the locations of each display associated trigger zones defining the area in which the mobile application sends a request to the display [Dav+14]. Of course, with the request for personalised content on certain displays viewers also reveal their location – and it is possible to create a comprehensive dataset of user-related location traces. The authors, however, draw on pre-existing trust relationships between viewers and content providers where the viewer's location is only shared with the personalisable applications that have been activated.

Previous work in the context of mobile phone based interactions with public displays has also identified other privacy-related concerns of viewers. For example, in the work of Cheverst et al. [Che+05], users were able to use mobile phones in order to share and retrieve pictures on a photo display using Bluetooth. Users were particularly concerned regarding the accessibility of the pictures that have been shared with the display and highlighted the importance of providing full control over the data (e.g. to be able to remove shared data).

2.3.4.5 Comparing Interaction Modalities

Different interaction modalities can have varying impact on the viewer perception and user experience of public displays. Researchers have used analytics techniques to compare the effectiveness of displays and content depending on the interaction modality that was supported by the display. Alt et al. performed the "Waiting Room Experiment" in which the authors deployed a public display equipped with both touch and gesture sensors and the "Soap Bubble" application that supported interactions through sensors [Alt+13]. The authors tried to understand if interacting with a display has a cognitive effect on the viewers' memory to remember the content better compared to non-interactive displays. The authors made heavy use of the available sensing technology and collected the entire depth video stream and performed interactions categorised by the interaction type. To measure the actual effectiveness (in this case, in the form of content memorability), the authors conducted a survey at the end of the study to understand whether viewers remembered certain content based on the interaction technique that was used. Parra, Klerkx, and Duval have conducted an experiment in which they compared two different strategies to communicate the interactivity of a display to the viewers: silhouettes and mirrors of passers-by transferred to the public display in real time [PKD14]. In particular, the researchers utilised a combination of visual computing techniques and interaction logs were utilised to quantify and describe viewer interactions in the form of a flow diagram for each of the two interaction modalities. Similar to other deployments, the collected video analytics and interaction data is processed in real time and used to change the content on screens and automatically react to viewers approaching the display [PKD14].

Generally, viewers can *interact* or *engage* with a public display in many ways. She et al. have performed a literature survey in which they found a number of interaction modalities – including Bluetooth-based sensing, Near Field Communication, video processing or direct touch [She+14]. The heterogeneity of interaction modalities requires different data collection and reporting techniques to support digital signage analytics and the automated measurement of user engagement.

2.4 Reporting

The majority of the work in digital signage analytics has focused around the development of novel models for describing user behaviour (Section 2.2), and the collection of sign and viewer interaction events (Section 2.3). In this section, we will introduce work that mainly describes the reporting of analytical metrics and insights in the context of digital signage, and how such systems have emerged from Web analytics. In particular, we firstly introduce systems that were developed to report purely systems and content focussed insights in digital signage. Secondly, we show how modern analytics about viewer behaviour has emerged from Web analytics, and provide an overview of visualisation techniques used in the context of digital signage analytics – including systems that have been developed as part of an overall product, and systems and approaches that are part of research projects.

2.4.1 Statistical Reports about Systems and Content

In Sections 2.2 and 2.3, we have introduced a set of commercial products and research work specifically tailored toward the signage analytics domain, some of which capture metrics similar to these defined for Web analytics to measure and report user interactions and engagement with displays. Whilst many of the presented research papers focus on the extraction and collection of certain metrics, commercial products typically also feature the reporting of collected insights and relevant analytics data.

Digital signage analytics traditionally focus on the reporting of logs captured on the display and specifically focusing around metrics captured about the display itself [CAY; Loo17]. CAYIN Super Reporter and Look Digital Signage are examples for products that provides administrators with the ability to retrieve statistical reports about the log of played content, and aggregations of the duration content was played back across the display network [CAY; Loo17]. In addition, Look Digital Signage includes a dashboard that provides real-time insights into the current health status of connected displays and the associated signage player software, e.g. to identify displays with a hardware or software problem and such that require maintenance. Reports are composed of simple pie and bar charts, with an interface allowing administrators to filter for specific displays and content. CAYIN advertises their system as a tool for billing purposes to compute the durations a specific advert has played across all displays part of the network. While the aggregation and report generation capabilities are limited, the tool features the exporting of reports into third-party applications (e.g. spreadsheet calculation tools) for further analysis and aggregation. Similarly, Look Digital Signage includes capabilities to create statistical reports of content played across all connected displays with the particular focus on providing a “proof-of-play” for further use including billing purposes [Loo17]. Typically, such proof-of-play reports include a “detailed report confirming each and every time a campaign has been played back” [OnS17].

Commercial display analytics systems that focus on the reporting of audience numbers and demographics (introduced in Section 2.3.2, p. 19) also consider viewer privacy. For example, Intel AIM [Int18] provide anonymised insights of viewers (e.g. report aggregated numbers of specific demographics and age ranges only). Additionally, video streams are analysed on or close to the sensor neglecting the need for streaming and storing video recordings of the audience.

2.4.2 Visualisations of Analytics Data

Different visualisation techniques have been used both in the research domain and by commercial products. This section focuses on specific visualisation techniques including *funnel and flow diagrams*, and *heatmaps* used to describe viewer behaviour and movement patterns in the context of digital signage.

2.4.2.1 Funnel and Flow Diagrams

Researchers have used various techniques to report data pertaining to signage analytics and user interaction. Parra, Klerkx, and Duval use a flow diagram to show aggregates of different interaction stages of passers-by (walking by, noticing the display, triggering content, and actively interacting) with a display [PKD14]. The flow visualisation can be best compared with modern features provided by modern Web analytics frameworks such as Google Analytics in which user behaviour and sequential Web site visits are reported using funnel visualisations that show the proportions of users accessing certain sites in order [Goo18f]. For example, knowing at which interaction stage most of the users have stopped interacting with the display could reveal potential problems with the signage player or interaction modalities. Researchers have developed a number of audience description models (initially described in Section 2.2), some of which consists of elements about the description of audience behaviour [MM11; PKD14]. As part of the Audience Funnel Framework, Michelis and Müller have developed solutions to describe user behaviour and navigation patterns in the vicinity of displays in which they utilise funnel-based expressions and visualisations. Specifically, Michelis and Müller describe the conversion rates of viewers from across the four interaction phases (passing by, subtle interaction, direct interaction, and multiple interaction) giving a detailed understanding of which proportion of passers-by were retained throughout the phases. For example, as part of their study, the authors collected a set of interaction data and were able to measure a 33% conversion rate of viewers transitioning from passing-by to subtle interaction, and 95% from subtle interaction to direct interaction. In this concrete examples, the authors have chosen a mix of bar chart and flow diagram in which each bar represents the proportion of viewers compared to the number of passing-by, overlaid with the conversion rate of viewers between each phase [MM11]. Such insights allow developers and display providers to understand which phase transitions work effectively and which require attention and improvements. Of course, such reports are only applicable to display deployments that can be mapped to the original Audience Funnel Framework, i.e. provide a certain level of interaction capabilities. Parra, Klerkx, and Duval specifically showed how the dataset could be used to provide conversion rates (number of people interacting with the display relative to the total number of passers-by) per hour to understand at which times certain interaction stages are the most likely for a user to go through [PKD14]. Parra, Klerkx, and Duval have further chosen funnel visualisations to show the flow of viewers going through a number of phases – clearly indicating at which phase a proportion of viewers is transitioning to the next phase [PKD14].

2.4.2.2 Heatmaps

The use of heat maps can be particularly useful to report and visualise user navigation flows and behavioural patterns linked to locations for both in- and out-door environments. Girgensohn, Shipman, and Wilcox developed an activity recognition system that plots average speeds, people counts and behaviours onto a spatial map enabling the user to understand how proportions of visitors behave in their space [GSW08]. Whilst this system is not built specifically for signage analytics, locations of digital signs could be plotted onto the map as an additional piece of information to reveal more insights into potential influences of the signs on viewer behaviour. In a related approach, Williamson and Williamson make use of heat maps to visualise the flow of pedestrians around an out-door display based on different kinds of content displayed on the screen as viewers walk by [WW14]. The aim of such reports is to help understand researchers and content creators which how certain content and display locations influence on the viewers walking and behaviour patterns, which content draws more attention and, specifically in this case, which content leads to a change in behavioural and navigational patterns of the passers-by.

2.4.3 Analytics Reporting in Related Areas

The creation of analytics reports regarding the behaviour and interactions of individuals has been subject to prior work in areas closely related to digital signage such as Web analytics and retail.

2.4.3.1 Web Analytics

Web analytics form one of the foundations for the creation and use of user interaction and behaviour reports. Whilst a number of differing definitions for Web analytics exist as previously noted by Peterson [Pet04], the authors have adopted a more general definition highlighting the large number of potential data sources for Web analytics systems and the use of the insights created:

> "Web analytics is the assessment of a variety of data, including Web traffic, Web-based transactions, Web server performance, usability studies, user submitted information and related sources to help create a generalised understanding of the visitor experience online." – Peterson [Pet04]

As highlighted in the quote above, Web analytics considers a number of different datasets and data sources to create results. Jansen [Jan09] describe two main methods for capturing data relevant for analytics reporting: (1) Web server log files, and (2) page tagging. Web server log files are page access logs typically created by the Web server application on each page or content request – including additional meta data such as the IP address of the requesting client, cookies, and details on the Web browser and operating system of the client [Jan09]. Page tagging describes the method of using hidden JavaScript snippets that "send information about the page and the user back to a remote server" [Jan09]. This approach allows, for example, the capture of in-page user behaviour such as scrolling and is commonly used by modern analytics services.

Given the large amount of analytics-relevant data captured on Web servers, a number of insights can be gained through appropriate reports and researchers have developed techniques that allow the extraction of insightful and comprehensive information from such access logs and other data sources. Cooley, Mobasher, and Srivastava brought together a set of techniques for the extraction of patterns by analysing web transaction logs of users accessing websites, and are typically used by Web administrators and potential content creators [CMS97]. The authors have specifically explored the creation and use of association rules of users interactions across multiple Web sites. Such association rules enable the reporting of, for example, the proportions of users who accessed a specific Web page upon previously visiting a different page of the same domain, and the reporting of the proportion of users who entered a Web site through a specific page compared to other entry pages. Specific example reports developed by Cooley, Mobasher, and Srivastava further include those of the type: "*40% of users who accessed website A, also accessed website B*", "*20% of users who accessed website C, purchased product X*", "*30% of users who searched for keyword Z accessed website Y*", and "*30% of viewers of website A live on the East Coast*" [CMS97].

The overall goal of Web analytics is generally to improve the user experience in the Web based on the insights gained [CMS97]. According to Cooley, Mobasher, and Srivastava, the reporting of such user behaviour has the potential to enable targeted advertising and provide, at the time, novel and unique insights into the user behaviour and navigation patterns to inform a better and more effective website design – ultimately with the goal to improve the quality of the website [CMS99]. The quality of the reporting can be even improved with the identification of individual users and *user sessions* while "the goal of session identification is to divide the page access of each user into individual sessions" [CMS99]. This enables the creation of more detailed analytics reports in which the identified session is part of the aggregation. For example, whether the user's visit of *website Y*, after previously visiting *website X*, is counted as a consequential page visit or as an entry page to the site would depend on whether both visits have been identified as part of the same or a different session.

The primary application area for Web analytics is to explore and understand how users interact on Web sites [Jan09]. Such analytics reports can be used, for example, to measure and understand human social behaviour as in the work conducted by Gonçalves and Ramasco [GR08] who analysed access logs captured on the server side from a university website. Based on the log files, the authors developed access and usage patterns and derived various information such as sleep and work patterns of individuals [GR08]. Other domains include journalism in which Web analytics are used to closely monitor the success of news stories – directly influencing news stories to prominently feature on home pages in the form of a gatekeeper [Eds14].

The tracking of individuals across distinct Web sites can be a potential invasion of the user's privacy and has been identified as a particular issue [KW09]. We note that attempts at protecting user privacy exist. For example, commercial analytics services such as Google Analytics implement measures that allow owners of Web sites to activate the anonymisation of personal data (i.e. the visitor's IP address) in order to ensure reduce the invasion of privacy [Goo18d]. Once activated, all analytics events relating to a user are stored against a masked IP address only and cannot be tied to an individual anymore.

2.4.3.2 Retail

In the retail context, retail-specific datasets have been used for digital signage analytics as a way to evaluate the effectiveness and success of advertising campaigns, and therefore also public displays deployed in a retail context themselves [AMS12]. For example, to determine the level of success of an advertising campaign, retailers and advertising agencies analyse point-of-sales information that have been gathered over long periods of time and conduct surveys and lab-like observations of viewer responses to display content [Sla10a]. In the context of advertising on public displays, the effectiveness of a display can be partially defined through the effectiveness and success of an advertising campaign which is immediately linked to the increase in sales of the advertised product. To automate this process, Senior et al. have deployed digital signs and a face recognition system in a retail context to identify and track people – and enabling the use of such a system with already deployed surveillance cameras [Sen+07]. The authors explicitly identified customer counts and mention the measurement of "display effectiveness" as use cases for their system. The display effectiveness in this context has been based on the dwell time of viewers in the immediate vicinity of the display and their trajectories around the display inside a retail store. With an increase in the dwell and interaction time of viewers with a display, the display was described as more effective [Sen+07]. Of course, linking this information with trajectories beyond the vicinity of the display inside a shipping mall or similar environment could allow us to draw conclusions on the display effectiveness with regards to behaviour change.

The similar set of technologies was used to capture valid transactions and link such transactions to customers who have viewed a certain piece of content or advert prior to purchasing a product. Rai, Jonna, and Krishna conducted research into the detailed analysis of customer behaviour in the retail space – up to the level of detecting objects that customers have placed into their shopping card while still navigating through the store [RJK11]. Such a high level of tracking and insights about customers can be best compared with current Web analytics techniques in which user interactions, and placements of products into virtual shopping cards are recorded and assigned to individuals. Rai, Jonna, and Krishna elaborate on the emergence of novel business opportunities for retailers based on their system, including the delivery of location- and context-aware advertisements to public displays close to the visitor [RJK11]. Example content could include promotional offers based on the location of the visitor and the products that have been placed into the shopping cart. Pervasive computing technology and the ability to track visitors and purchases in the offline-world will likely change "the face of retail" [KSO11]. Krüger, Schöning, and Olivier describe the replacement of traditional, paper-based signage with digital displays as one of the potential changes in modern retail – providing a number of opportunities to improve the experience for customers by, for example, delivering personalised messages and recommendations [KSO11]. To provide visitors with additional information about products, Strohbach and Martin developed a platform for context-aware public displays specifically designed for the use in retail environments [SM11]. The system features a modular architecture in which a number of "context agents" can be plugged in – each of which representing a separate data source such as RFID sensors and visitor tracking to allow the display to automatically adapt to the visitor and their preferences. For example, the built-in RFID tracking recognises RFID-equipped products that a visitor places on the screen and automatically visualises additional information about the product such as descriptions or the price [SM11]. The ability to plug in sensors and therefore external data sources can be one way to feed data initially captured by analytics systems back into the digital sign to improve the user experience.

The systems described are characterised by their desire to collect large amounts of data about individuals in order to understand the effectiveness of public displays in influencing customer decisions (e.g. success of advertisement campaigns by linking customer behaviour before/after interacting with a display [RJK11]). The desire for capturing and processing large amounts of data relating to individuals, however, introduces a number of privacy-related implications as highlighted by Tene and Polonetsky [TP13]. The authors particularly argue for the need of a "legal model where the benefits of data for organisations and researchers are shared with individuals" [TP13] – providing some insights to users about the usage of their personal data.

2.5 Automated Use of Analytics Data

Analytical insights can be used to dynamically change the content on public displays and digital signage to ensure the content better suits the current contextual events, e.g. tailored toward the demographic of an audience. The common (commercial) use of such systems is the provision of advertisements targeted to a specific group of people (Section 2.5.1) and support for the creation and scheduling of content (Section 2.5.2).

2.5.1 Targeted Advertising

The use of analytics data for the placement and messaging of targeted advertisements to individuals has originated in the Web domain and has emerged from efforts to create personalised Web experiences through the use of data mining on Web usage logs [MCS00]. Google AdSense [Goo17] is an example of a commercial product that mines user interests and analysis the Web site on which the administrator has placed the advert. The resulting advertisement is highly context-aware and, at the same time, tailored toward the viewer's interests. In contrast to digital signage analytics, such modern Web advertising tools track users across multiple websites and domains, and therefore provided insights into users' traces throughout the Web. As a result, companies gain a detailed understanding of the user behaviour and individuals are likely to receive a different personalised advertisement even though they are visiting the same Web site.

Similar to the trends in Web analytics, researchers are working on systems that support targeted advertisements "in the real world" for digital signage and public displays. While the use of actuated displays and the support for the delivery of personalised and targeted content through public displays and digital signage were referred to in the context of improving the user experience [Sla16], such systems are used specifically for targeted advertising. For example, the previously described Proximity Toolkit developed by Marquardt et al. [Mar+11] has been used as a foundation to enable researchers the development of novel interactive content including "attention-demanding advertisements" [Mar+11]. The purpose of the system is to support advertisements that dynamically adapt to the attention level and location of the passers-by in the vicinity of the display to grab and retain the passers-by attention. The application relies on real-time information computed and provided by the underlying analytics system of the toolkit based on the Proxemic Interactions model developed by Ballendat, Marquardt, and Greenberg [BMG10] and described in more detail in Section 2.2 (Audience Models and Metrics).

While the Proximity Toolkit uses spatial and contextual information about the attention level captured in the vicinity of the display to support the delivery of tailored advertisements, other systems were developed that capture more detailed insights about the audience and individuals. The use of demographic information for the actuation of displays and the selection of *targeted* content on digital signage is a particular focus of commercial systems for enabling *targeted advertising* to individuals passing by displays [Far+14; Sla10a; Tia+12]. To enable the development and implementation of such applications (i.e. dynamic display and content changes), commercial providers utilise visual analytics techniques and frameworks that feature appropriate APIs specifically designed for the purpose of creating dynamic and interactive content. SCALA provide a visual analytics tool that supports the “plug-in of sensors” [Sca] and was designed to “adjust automatic messaging to reflect customer habits and trends” [Sca] – practically providing ways for content creators to build sensors that are connected to the analytics systems and integrated as part of public display applications. Other commercial suppliers have worked on applications that make analytics insights available instantly via a network service [Qui16] to grab the attention of a passer-by and make viewers aware of interaction capabilities of kiosks [Sla16]. As part of their product line, Meridian Kiosks offer an analytics framework and corresponding digital signage applications and signage players that enable “more targeted messaging and advertising” [Sla16] with the goal to “incentivise people to touch the screen and engage” [Sla16]. In such cases, receive a control feed that enables the distribution of targeted messages and adverts to individuals based on their behaviour in the vicinity of the display, and can also be used to inform the quality and effectiveness of certain kinds of content [Sla16].

To better tailor advertisements to individuals, Tian et al. developed a system that orchestrates insights from a number of on-screen and external data sources including video analytics and sales statistics [Tia+12]. The developed “intelligent advertising framework” consists of a learning module that captures how well demographic groups responded to certain ads (e.g. measured by an increased number of views), and starts delivering out ads to groups of people of the same demographics. External data sources, in particular sales statistics from point-of-sales terminals, are used to capture the effectiveness of such advertisements in real time. Display owners and advertisers were able to additionally define rules for content to be shown preferably to a certain demographic group – in this case, combinations of age ranges and gender. The resulting system was unfortunately not extensively evaluated though, according to the authors, attracted significant attention from the advertising industry [Tia+12]. NEC offer a similar system that enables “real-time marketing” [NEC13] through an integrated digital signage solution. The system utilises insights captured through analytics (e.g. point-of-sales statistics) to drive the content shown on displays [NEC13].

While the majority of systems use “anonymous video analytics” (initially described in Section 2.3.2) for enabling targeted advertisements, Farinella et al. went one step further and support face reidentification for the selection of targeted advertisements [Far+14]. The authors emphasise the value in being able to determine if a viewer present in front of the display has seen a particular advert previously, which would highly impact the selection of an advert [Far+14]. Of course, such an approach is not anonymous and does indeed collect unique information about a visitor enabling the system to potentially build up deep insights such as navigation and visiting patterns about an individual.

Langheinrich and Schaub [LS18] highlight privacy-related risks that emerge with the increasing amount of “custom-tailored and context-aware services” [LS18]. The authors discuss a set of approaches of system designs for pervasive computing to help improve user privacy whilst still enabling the provision of tailored services for individuals.

2.5.2 Content Scheduling and Content Creation

Displays and deployments that adapt to local analytics and contextual events are often referred to as 'situated displays'. Research has been conducted to explore the use of (real-time) analytics data for the selection of content shown on displays (e.g. to advertise technical capabilities of the display to the viewer passing by [Dav+14]). A set of previously presented papers in Section 2.3 (Data Capture) provides an "attention grabbing" feature that uses real-time analytical insights to capture and react to passers-by with the purpose of making them aware of the interaction capabilities of the public display and ultimately to retain their attention for a longer duration. For example, the public display system developed by Parra, Klerkx, and Duval [PKD14] uses real-time analytical insights that are fed back into the digital sign to automatically select appropriate content to be shown to individuals in vicinity of the display. Similar to the Looking Glass [Mül+12], the system developed by Parra, Klerkx, and Duval shows a "digital mirror" of each passer-by that dynamically follows the person as they move around in the vicinity of the display – with the aim of drawing in more attention from the audience and hence improving the visibility of the display [PKD14]. As part of the interaction models established by Vogel and Balakrishnan [VB04], the authors have developed a prototype system that utilises the viewer's location adjacent to the display to help transitioning across interaction zones by automatically adapting the content and activating the appropriate interaction modality [VB04], as described in Section 2.3.3 (Audience Engagement and Movement). Insights about viewers such as their demographics are further used to influence the kinds of content that are displayed in real-time. Whilst we have introduced a subset of such applications in Targeted Advertising (Section 2.5.1), researchers worked on systems that go one step further and tailor content to individuals recognised and identified by the system – across multiple display locations. This includes systems developed by Gillian et al. [Gil+14] and Farinella et al. [Far+14] that include built-in face recognition to identify viewers, and deliver highly tailored and personalised content based on their preferences such as personal calendar information.

Analytics data has also been used to inform the creation of new kinds of content and drive the development of novel interaction and gesture modalities for digital signage. Examples of such systems include commercial products that capture audience numbers and demographics (described in detail in Section 2.3.2). Slawsky describe such commercial tools as a way to allow content creators and administrators measure the effectiveness of their content (in this case effectiveness is measured by the number of views a piece of content has received by the audience), allowing content creators to understand which kinds of content produce better numbers and are therefore more effective [Sla11]. Researchers have used analytical insights including those collected through observational studies to inform the design of novel content and interacting gestures for public displays [HRD11]. In this work, Hardy, Rukzio, and Davies captured attention levels per content and interaction type which enabled the authors to extract and determine which kinds of content achieved the desired results (i.e. more attention and interactions). Ultimately, such insights allowed the capture of preferred gesture types enabling researchers to focus on the development and deployment of such gestures that perceived more interactions (e.g. due to their simplicity) [HRD11].

2.6 Analysis

The related work described spans across foundational research on audience interaction models for describing viewer behaviour in the immediate vicinity of a display, data capture and reporting techniques such as video analytics, and work that automatically utilises analytics data to adapt novel forms of interaction modalities. In this section, we provide an overview of the coverage of the body of related work regarding the three areas of challenge (analytics data, reporting and the automated use of analytics data), and subsequently provide an analysis of the suitability of the work for use in the context of open pervasive display networks by identifying and highlighting limitations for each of the characteristics (openness, pervasiveness, and networked).

2.6.1 Evolution and Coverage

In order to allow us to highlight the coverage of the body of related work and the changing research focus over time, we first categorised each piece of work into one of the three areas of challenge. To provide a better understanding of our categorisation, we defined each of the areas as follows:

Analytics Data Capture Systems, probes or solutions that are related to capturing analytical insights about viewers / audience, users of a system, and the digital sign itself. Additionally, we consider work regarding the development of novel capture techniques such as visual computing algorithms.

Reporting Systems, probes or solutions that specifically focus on the creation and dissemination of relevant analytics reports, including work that focuses on the presentation of relevant or novel metrics.

Automated Use of Analytics Data Systems, probes or solutions that utilise analytics insights in real-time to dynamically adapt the content shown or the available interaction modalities to viewers present in the vicinity of the display. Additionally, we include work on retrospectively informing the design of (interactive) content or display deployments based on analytics.

Due to the large amount of related work that was specifically conducted on investigating and enabling novel interaction modalities (e.g. via gaze), we additionally consider **interaction** as a fourth area – allowing us to provide further insight into the research focus of the related work. Of course, novel interaction modalities can be used as a driver for capturing relevant analytics data (i.e. viewer interactions and engagement with the display). However, we assign

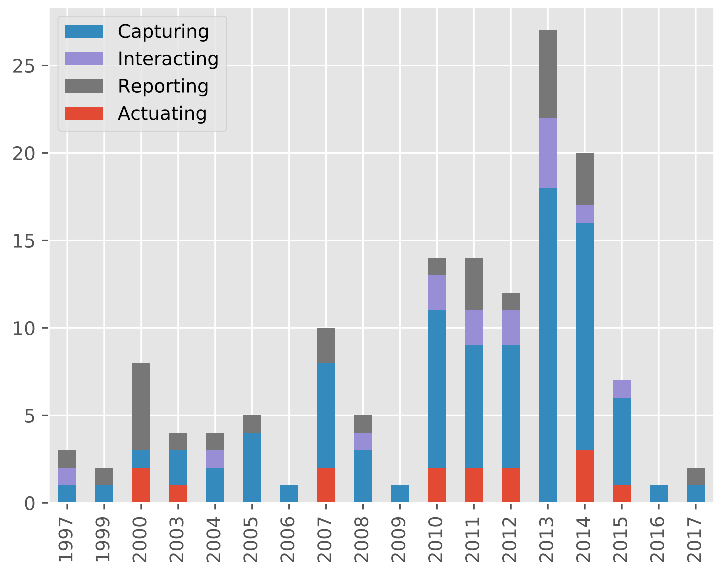

Figure 2.3: Distribution of related and relevant digital signage and Web analytics work by their publication date grouped by the main focus of the paper into four categories: *capturing* of analytics data, *interacting* of users, *reporting* of relevant analytics insights and the use of analytics for *actuation*.

related work to the 'interaction' category only if it was not specifically focusing on enabling analytics.

Figure 2.3 shows the coverage of the related work for each year as a stacked bar chart (due to the scope of some of the introduced research papers, a subset of papers may be assigned to multiple categories at the same time). We observe that the majority of the work focuses on areas of data capture and interaction – i.e. enabling the capture of novel interaction events from the audience. Little work has been carried out on the creation of analytical insights and reports, or that utilises such reports for the actuation of digital signs. We note that the increased trend of developing novel forms of interaction modalities – likely to be grounded in the opportunities that have emerged with technical advancements. For example, gaze [SHT10], gesture [Gil+14] and mobile phones emerged as new ways for viewers to interact with displays in recent years. Additionally, we observe a constant tendency to focus on analytics data capture techniques (e.g. user mobility tracking in spaces [Jay+16; WW14] or the detection of attention states of viewers [WBG12]). The common thread across the related work lies in the focus on single or small scale display deployments with analytical insights typically provided about the audience (e.g. demographics and age ranges) and their behaviour (e.g. interaction phases). For example, the introduced audience models and metrics (Section 2.2, p. 14) define a detailed language and understanding of how viewers navigate and behave in the immediate vicinity of a display (e.g. "passing by" and "viewing and reacting" [MM11]), perform certain forms of interactions (e.g. implicit and subtle interaction [VB04], and direct interactions [MM11]) and the consideration of potential follow-up interactions (e.g. "follow-up actions" [MM11] or "conatation" [She+14]).

In the following section, we will address the shortcomings and limitations of related work in the context of open pervasive display networks.

2.6.2 Suitability for Pervasive Display Networks

Our work is motivated by the shift from closed networks of displays towards *open pervasive display networks*, drawing on the scenario introduced in Section 1.2 (p. 3). Specifically, such networks of displays feature a unique set of characteristics regarding their 'openness', 'pervasiveness' and 'networked' aspects. Considering the existing body of work in digital signage analytics in terms of these characterisations highlights a set of limitations and shortcomings.

Openness

With pervasive display networks becoming more *open*, the number of potential stakeholders who contribute to the display network significantly increases compared to closed display networks. Previous work has identified such stakeholders to include display owners, content providers, space owners and specifically also include viewers as a fourth stakeholder [VO11; AMS12; Cli+14]. Each of these stakeholder groups are likely to be composed of a high number of individual stakeholders.

The existing body of work, however, typically focusses on capturing and reporting analytics for single stakeholder entities only. As an example for the body of work that focuses on data capture and reporting, the system-based monitoring tools introduced are predominately designed to provide display owners with an overview of the entire network of deployed digital signs (e.g. Esprida [Esp17] and CAYIN [CAY]). More sophisticated analytics tools (e.g. Intel AVA [Cav11] and Fraunhofer SHORE [Frab]) additionally provide insights into the audience demographics (including age ranges and gender using cameras mounted on displays). Such audience analytics reports, however, are targeted for display owners as the single stakeholder entity. A similar trend can be observed in work on feeding (analytics) data back into the sign, e.g. to support novel forms of interactions. For example, systems that utilise insights on viewer demographics [Tia+12], audience attention levels and locations [MG12] and interactions with physical objects in the vicinity of the display [Sla16] have been designed for use with either a single display deployments (e.g. proximity-based interactions [MG15]) or in the context of environments that are controlled by a single stakeholder entity (e.g. [Tia+12]). However, with displays becoming open ecosystems to which a range of stakeholders can contribute, the complexity increases in accommodating requirements and constraints imposed by individual stakeholders and the high number of potential contextual changes to be fed into the sign. It will likely become highly challenging to, for example, create a definite display content schedule that accommodates all these requirements and is capable of dynamically responding to analytics events in the vicinity of a display.

We were unable to identify prior work that considers *openness* as a characteristic across the data capture, reporting and automated use of analytics data areas of challenge. In particular, work that considers capturing and combine relevant datasets across a range of stakeholders (also including viewers as an equal stakeholder) leveraging the potentially rich ecosystems through which viewers navigate and in which viewers interact. Drawing on the introductory scenario, analytics systems that consider events from multiple data points, likely to be owned by multiple stakeholders, will be required in future *open* display networks. Additionally, such systems will need to include mechanisms for resolving potentially conflicting requirements and constraints when feeding analytical insights from various stakeholders back into the sign.

Networked

Existing work focussed on capturing and reporting analytics regarding individual displays or networks owned by a single stakeholder entity. In particular, the techniques described focus on capturing and analysing data from the perspective of a display, i.e. provide insights about displays and their audience. For example, commercial solutions capture a set of system-related information from displays that are part of a single network such as logs of content played and the state of display players [CAY; Rem17; OnS17]. Additionally, more advanced solutions feature visual analytics techniques to capture (anonymised) audience numbers and demographics from individual displays including gender, age ranges and in some cases even the mood of members of viewers [Int18; IBM13; Tia+12; Alt+12].

In contrast, with the emergence of interconnected networks of displays we will be able to observe potentially complex and spatially distributed viewer interactions and engagements with displays and the content shown that will be difficult to describe with conventional analytics techniques. Instead, a shift towards a *viewer-centric perspective* on analytics will be required to enable the creation of reports that describe how viewers experience digital signs and content when moving across spaces. For example, instead of capturing and reporting the average demographic group, analytics may provide further insights on the content viewers have previously seen, and the potential impact of displays on their behaviour and movement patterns – almost providing an equivalent to a click-through event given in Web analytics. We were unable to identify prior work that focussed on providing a viewer-centric perspective in the context of open pervasive display networks.

Pervasiveness

The final characteristic of future display networks is the *pervasiveness*: public displays appear embedded into urban environments and are becoming omnipresent to the viewer. With displays becoming more pervasive and ubiquitous, the amounts of data, data stakeholders and users is likely to grow consistently – creating a need for novel analytics systems that facilitate the requirements emerging from the unique characteristics from future open display networks. Such novel analytics systems, however, impose risk of revealing highly sensitive and potentially privacy-invasive insights about individuals. For example, researchers have focussed on developing accurate indoor location tracking technologies that work within a limited spatial area, e.g. by utilising Wi-Fi hotspots [GJC04; MB13; Air13] with the overall goal to gain a better understanding about the behaviour and navigational patterns of individuals. Some work considered to additionally capture the viewer's identity to provide certain personalised services such as access to calendars [Gil+14], or proposes the knowledge of viewer identities to enable the association of distinct events (e.g. display views and purchases) [RS13].

However, in the context of open display networks with viewers becoming a participating stakeholder, we believe it to be crucial that analytics can be captured in a privacy-preserving way yet providing detailed viewer-centric insights into the viewer experience. We were unable to identify work that supported the capture, reporting and subsequent automated use of analytics insights that was designed in a privacy-preserving way for the use in open display networks or pervasive computing environments.

2.7 Summary

In this chapter we provided a detailed overview of research efforts and state-of-the-art systems in digital signage analytics and related domains. We introduced audience models and metrics describing the requirements and opportunities for digital signage analytics specifically regarding audience behaviour in the vicinity of displays. In the context of data capture, we described early work on capturing analytics-relevant data on public displays – predominantly focused around the performance and reliability of systems (i.e. individual display nodes), and more recent work on using visual analytics techniques to capture audience counts and demographics, interactions and behaviour patterns in front of the display. Systems for reporting of analytical insights include statistical reports about the system performance and the content displayed across a network of displays (e.g. to provide a "proof of play" for billing purposes), and visualisation techniques used to express analytical insights pertaining to public displays. In terms of automated use of analytics data, we described systems that support the delivery of targeted advertisements based on audience demographics, and those that utilise more sophisticated insights (e.g. the viewer's identity) to schedule personalised content in real time to a display. Finally, we gave an overview of analytics in related domains including user tracking across Web sites, and capturing and providing analytics in a retail context. An analysis of this body of related work in the context of open pervasive display networks revealed that analytics are typically considered for displays owned by single stakeholder entities – whilst in *open* networks of displays a large amount of stakeholders contributes to the display deployment. Additionally, the body of work imposes limitations to the extent to which analytics is captured and reported on the viewer-centric experience with a display network across multiple locations – whilst preserving the viewers' privacy.

Chapter 3

Analytics Data Capture and Generation

3.1 Overview

In this chapter, we investigate the identification and collection of relevant analytics datasets. We begin by exploring the challenges and opportunities that openness and pervasiveness bring to the capture and generation of digital sign analytics and present a framework for analytics data sharing. A key finding from this work is the providing viewer-centric analytics requires the combination of conventional sign analytics and user mobility tracking data (see Figure 3.1). Hence, we provide an architecture for an analytics platform for the collection and processing of traditional and display-oriented analytics data, and describe two different approaches for the collection of viewer mobility and interaction logs: using client-based tracking technology through the viewers' mobile phones and, in contrast, the use of infrastructure-based tracking through Wi-Fi fingerprinting. Finally, we describe the use of mobility models for the creation of synthetic analytics as an approach that preserves viewer privacy and, simultaneously,

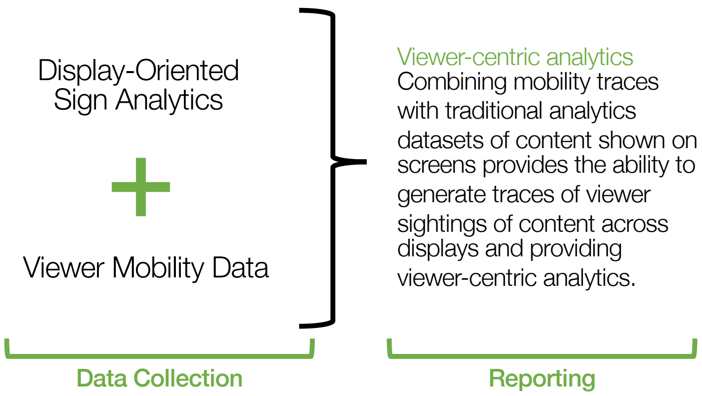

Figure 3.1: Fundamentals of the generation and collection of viewer-centric analytics in digital signage.

provides us with the foundation to generate signage analytics reports that have previously not been possible due to the lack of comprehensive viewer tracking.

Excerpts of this chapter have been published in the following peer-reviewed publications:

1. Mateusz Mikusz, Sarah Clinch, Rachel Jones, Michael Harding, Christopher Winstanley, and Nigel Davies. "Repurposing Web Analytics to Support the IoT". in: *Computer* 48.9 (Sept. 2015), pp. 42–49. ISSN: 0018-9162. DOI: 10.1109/MC.2015.260. URL: http://doi.org/10.1109/MC.2015.260

2. Mateusz Mikusz, Anastasios Noulas, Nigel Davies, Sarah Clinch, and Adrian Friday. "Next Generation Physical Analytics for Digital Signage". In: *Proceedings of the 3rd International on Workshop on Physical Analytics*. WPA '16. Singapore, Singapore: ACM, 2016, pp. 19–24. ISBN: 978-1-4503-4328-2. DOI: 10.1145/2935651.2935658. URL: http://doi.acm.org/10.1145/2935651.2935658

3. Mateusz Mikusz, Sarah Clinch, and Nigel Davies. "Design Considerations for Multi-stakeholder Display Analytics". In: *Proceedings of the 6th ACM International Symposium on Pervasive Displays*. PerDis '17. Lugano, Switzerland: ACM, 2017, 18:1–18:10. ISBN: 978-1-4503-5045-7. DOI: 10.1145/3078810.3078830. URL: http://doi.acm.org/10.1145/3078810.3078830

4. Mateusz Mikusz, Peter Shaw, Nigel Davies, Sarah Clinch, Ludwig Trotter, Ivan Elhart, Marc Langheinrich, and Adrian Friday. "Experiences of Mobile Personalisation of Pervasive Displays". In: *ACM Transactions on Computer-Human Interaction – TOCHI (in preparation)* (2018)

5. Mateusz Mikusz, Kenny Tsu Wei Choo, Rajesh Krishna Balan, Nigel Davies, and Youngki Lee. "New Challenges in Saturated Displays Environments". In: *IEEE Pervasive Computing* (2018)

3.2 Framework for Multi-Stakeholder Analytics Data Sharing

The analytics landscape in the context of open displays networks [Dav+12] is composed of a large set of stakeholders such as display and space owners, and content providers. The increasing number of public displays deployed in the environment and stakeholders involved lead to a complex network and ecosystem in which large amounts of analytics data can be captured regarding various aspects of the display network. We conducted an analysis of related literature to identify types of analytics data that can be captured, and subsequently designed a framework for the categorisation of data sources in order to identify relevant datasets that can serve as a foundation for the creation of analytics reports and oversee and understand important aspects of large open display networks. We believe that the categorisation of data types and data sources is an important first step to enabling the identification of opportunities for novel analytics insights to be gained in digital signage and public display ecosystems.

3.2.1 Stakeholders of Open Display Networks

A set of distinct stakeholders are involved in typical public display networks. In particular, related work in the field of open display networks and digital signage have identified a core set of four stakeholder groups: *display owners*, *space owners*, *content providers* and *viewers* [VO11; AMS12; Cli+14]. In the context of public display analytics, each stakeholder group is likely to be capable of collecting a unique set of data required for the creation of reports. The identified stakeholder groups are defined as follows.

Display Owners. Stakeholders owning a single display or network of displays are considered *display owners*. Typically, display owners have access to both the hard- and software that is powering displays and have an influence on the kinds of content and applications that are displayed. Examples for display owners include large commercial display network providers such as LinkNYC[1] or smaller research deployments such as UbiDisplays in which the displays are owned by the responsible research group [VO11].

Space Owners. Stakeholders who are in control of the physical space in which a display is deployed are referred to as *space owners*. For example, a display that belongs to a commercial advertising network might be deployed in a shopping mall that is owned by a distinct company.

Content Providers. An entity or organisation that is responsible for creating and supplying content through single displays or a network of public displays is described as *content provider*. Examples of content providers have been identified by Alt, Müller, and Schmidt as "event organisers, some types of service providers, and even passers-by" [AMS12].

Viewers. The 'targets' of public display networks and consumers of content and applications distributed through displays are referred to as *viewers* – who become stakeholders of the display network themselves. For example, viewers will be able to express preferences and contribute their own content to displays and therefore have a high impact and influence on the content displayed [Dav+12; AMS12].

In large-scaled open display networks each of these stakeholder groups can be composed of a large number of individual entities and organisations. Additionally, a single organisation might have multiple roles in a display network and therefore be part of multiple stakeholder groups at the same time. For example, in a shopping centre scenario, the organisation owning the physical space of the shopping mall might also own the display network deployed in the space that provides visitors with interactive way finding. Displays located inside the shopping mall used for advertising purposes, however, would show external content created and provided by stakeholder groups distinct from the display and space owner.

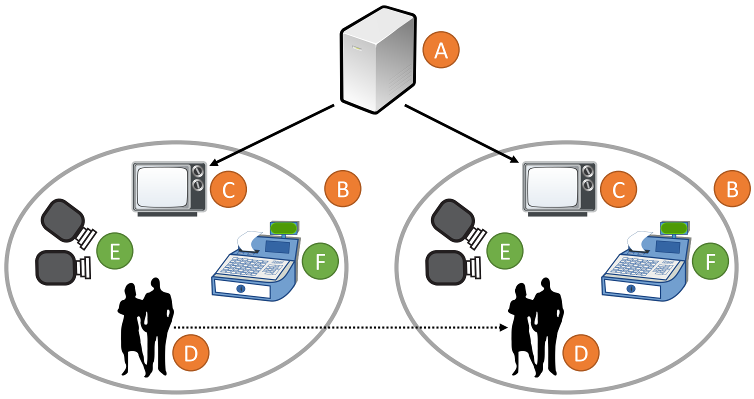

Figure 3.2: Stakeholders and their relationships in an open display network environment including content providers (A), space owners (B), public display owners (C), and viewers (D) who may move across multiple spaces. Example data sources include CCTV cameras deployed in a space (E) and point-of-sales terminals (F) in a retail context (first published in [MCD17]).

3.2.2 Stakeholder Data Collection Analysis

The motivating scenario described in Section 1.2 (The Need for New Forms of Signage Analytics) envisioned the stakeholders' ability to trace display sightings of passers-by with purchasing and interaction events beyond the vicinity of individual displays (i.e. seeing an advert for a soda drink and purchasing the advertised product at the nearest shop due to the display content shown). To support the analytics required for such deep understanding of cause and effects as part of the scenario, datasets from a large number of stakeholders need to be captured. As emphasised in Figure 3.2, individual stakeholders of a display network have unique opportunities for data collection—each required to support the opening scenario. In particular, Figure 3.2 shows the example of two spatially distinct display networks and a viewer group that is able to move between the spatial boundaries of one space into the other. Third-parties (Figure 3.2, A) such as independent developers, news agencies and advertising companies provide the content to the display network. Space owners (Figure 3.2, B) have access to a rich set of datasets relevant for analytics purposes. In the example of a retail space, the datasets can include closed-circuit television (Figure 3.2, E) and point-of-sales statistics (Figure 3.2, F). Displays are typically located within a space (Figure 3.2, C) but can be owned by a separate organisation. Viewers in display networks (Figure 3.2, D) become an important stakeholder and act as a connection point by moving across multiple spatially separated spaces.

Whilst some stakeholders might have access to overlapping datasets (e.g. audience demographics captured by display owners and space owners simultaneously), each stakeholder group have also access to a unique set of data and is able to collect relevant data to create analytical insights that can contribute to the generation of analytics reports for an entire display network. To better understand the kinds of data individual stakeholder groups are capable of collecting, and analyse the opportunities for the creation of novel insights by combining

[1]https://www.link.nyc

data from individual stakeholders, we conducted a literature review revealing the types of datasets relevant for digital signage analytics that have been considered by previous work. We categorised potential datasets and insights into six groups (Table 3.1): anonymous counting (i.e. anonymous statistics about a present audience), anonymous tracking (i.e. anonymous insights into the behaviour of an audience and their potential engagement with a public display), gesture recognition (i.e. datasets yielding insights into interactions and gestures performed by viewers), behaviour analysis (i.e. insights into audience behaviour within the vicinity of a display), pseudo-anonymous tracking (i.e. insights into audience behaviour across display locations) and contextual events (i.e. events relevant to the specific stakeholder such as sales statistics). The analysis of data collection opportunities for each stakeholder group is as follows.

3.2.2.1 Display Owners

Display owners are typically able to access and capture data via the presentation software directly on the digital sign. This includes a log of played content and content change patterns as well as potential failure states of the sign [Mik+15], and events of personalised content delivery to specific viewers based on their proximity to the display [BMG10; VB04; WBG12]. If the display is equipped with audience tracking sensors such as video cameras, the display owner may have access to further audience statistics and demographics [Int18], the orientation of viewers to the display [VB04], and their dwell times in the immediate vicinity of the display [Fraa]. Display owners can capture more comprehensive insights about the audience if displays are equipped with gesture recognition sensors (e.g. determining levels of interactions through gesture logs [PKD14]), and through face identification on displays equipped with appropriate video sensors [Far+14]. In particular face identification provides very detailed insights into the interaction and navigation patterns of viewers [Far+14].

3.2.2.2 Space Owners

Space owners have a broader view of the surrounding environment and are able to capture more comprehensive insights into the behaviour and navigation patterns of people who visit the space. This can include a general counting of people who are present in a particular area (e.g. in an area in which a display was deployed) [Mit+10], the time certain individuals have spent in a specified area [Sen+07], and simple entrance and exit counts of buildings and other spaces [Sen+07]. More sophisticated video analytics connected to a potentially already existing closed-circuit television can even extract the physical characteristics of people present in a space including their hair types, eye-wear, and colour of their clothing [Ham+09; TJS10]. Further insights can be gained by analysing the direction and speed of people in a space [WW14]—even across multiple cameras and locations [Yan+07]. Similar levels of insights can be achieved through Wi-Fi fingerprinting and location tracking achieved purely based on infrastructure sensors [Air13]. Specific examples of data collection opportunities from retail environments include the collection of sales statistics through point-of-sales terminals [Ven+07], conversion rates of purchases compared to the number of people present in a space [Sen+07], and additional point-of-sales statistics such as refunds and purchase cancellations [Ven+07].

3.2.2.3 Content Providers

Naturally, content providers have detailed knowledge about their application and the content that was supplied to a display network. This may well differ from the knowledge that is available to other stakeholders. For example, in the case of dynamic content such as an application that provides news feeds to displays, display owners might know that the news application was scheduled onto the display at a particular point of time, but they might not necessarily know the particular content that the news application distributed to the display. In this case, content providers would be able to capture the exact content that was created as well as the display that requested content [Mik+15].

In particular for displays and applications that support a user interactions, content providers are able to capture gesture interactions [HRD11], and user interaction events on the screen itself, e.g. through a touch-screen [Jac+10; VB04; Alt+11b]. Based on the user interaction logs, content providers are able to collect and create detailed navigation patterns within the display application—similar to insights that can be created by Web analytics tools [BS00] Whilst user interaction is one form of engagement, content providers can also capture engagement of users who follow up on content [Dob+11; Leh+12]. An example for an appropriate mechanism embedded in the content is the use of QR codes [GHN13].

3.2.2.4 Viewers

Viewers are able to capture and access a rich set of information and insights about themselves—such as their activities, location traces and life-logging. In particular, this includes their navigation and walking patterns [Ral+14] and additional information about the characteristics of their movements such as step counts [BH13] and dwelling durations [Ral+14]. With appropriate sensing technology such as smart glasses, viewers can are also capable of capturing gestures and gaze and therefore revealing, for example in a retail context, which products and displays have been noticed as viewers move through a space [Ral+14]. Further, modern smartphones that are typically carried around by a large proportion of viewers have sophisticated sensing technology integrated including accelerometers and GPS receivers [TJS10]. Viewers are also able to capture insights about other viewers in their vicinity, for example, to support indoor location tracking [LSC12].

3.2.3 Categorisation of Data Collection

Our literature review (summarised in Table 3.1) identified a wide variety of datasets that go beyond the scope of traditional signage analytics highlighting the fact that individual stakeholders have only limited insights and are unable to capture comprehensive insights about the behaviour patterns of the audience. For example, space owners are typically capable

Stakeholder	Anonymous Counting	Anonymous Tracking	Gesture Recognition	Behaviour Analysis	Pseudo-anonymous Tracking	Contextual Events
	Data Category					
Display Owner	Number of viewers facing a display [Gra+13; Int11; PKD14]; interacting with a display [Gra+13; Int11]; and viewer demographics (gender, age, etc.) [Fraa; Int11]	Orientation of viewers' heads [Jur+13; VB04; WBG12]; viewer distance from a display [BMG10; Jur+13; VB04; WBG12]; and viewer dwell time [Fraa; Jur+13; VB04]	Identification of pre-configured gestures [Gra+13; HRD11; Jur+13; VB04]; gesture logging [PKD14]		Re-identifying people across multiple screens [Far+14]	Content change patterns [Mik+15]; display state and failure monitoring [Mik+15]; presentation of personalised content based on proximity [BMG10; VB04; WBG12]
Space Owner	Number of potential viewers in display area [Mit+10; Sen+07; Ven+07]; time potential viewers spend in area [Sen+07]; and entrance and exit counts [Sen+07]	Physical characteristics of viewers (hair type, eyewear, clothing colour) through surveillance cameras [Ham+09; TJS10]		People tracking within video [Fraa; Mit+10; Sca]; re-identifying people [IBM13; Mit+10]; pedestrian models [Ant+06]; frequent visitors [GSW08]; fraud detection [Fan+09]; item interaction [Kra+05; Sca]	Direction, speed and location [WW14]; multi-camera tracking , iteYang2007; shopping cart tracking [RJK11]; WiFi fingerprints [Air13]	Number of sales [Sen+07; Ven+07]; conversion rate [Sen+07]; PoS events (refund, cancellation) [Ven+07]
Content Provider			Gestures for content selection [HRD11]; touch input [Jac+10; VB04]	Navigation patterns [BS00]; content interaction [Alt+11b]		Displays showing content [Mik+15]; QR code scans [GHN13]; user engagement [Dob+11; Leh+12]
Viewer				Viewer walking patterns [BH13; Ral+14]; step counts [BH13]; dwell times [Ral+14]; gaze tracking [Ral+14]; gesturing [Ral+14]	Data from interial sensors [TJS10]; indoor location with other viewers [LSC12]	

Table 3.1: Overview of analytics data that can be captured by stakeholders within open display networks, as grouped by category. Green cells indicate a rich set of data is held by a stakeholder in that category. Red cells indicates that stakeholders are typically unable to capture or hold information in the category (first published in [MCD17]).

Figure 3.3: Matrix composing the sources for data analytics collection (client- and server-side) and the dimensions of data sharing.

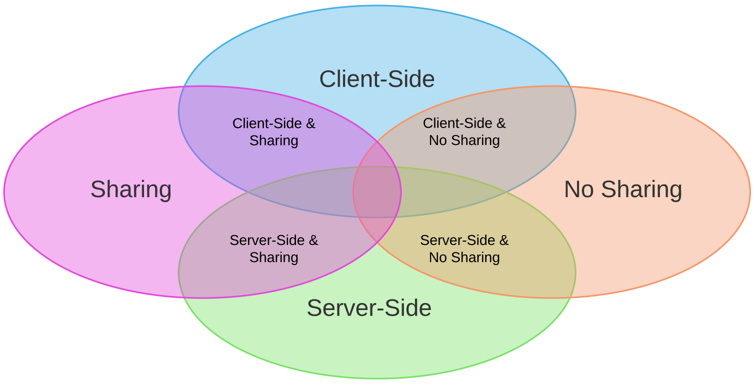

of capturing insights about the behaviour of visitors and retail statistics but have only little insight about the content shown on displays deployed within the boundaries of their physical space. Content providers know which displays have played and showed their content but have no or only very limited insights into the levels of interactions and engagement such as the number of views their single piece of content retrieved across a signage network. To allow stakeholders to gain more detailed insights into the effectiveness of their display deployments, applications and spaces, data will likely have to be merged and combined across multiple groups of stakeholders and organisations.

To help identify opportunities of combinations of datasets captured by distinct stakeholders that can lead to the capture and creation of novel insights, we designed a stakeholder data sharing matrix (Figure 3.3). This stakeholder sharing matrix considers the origin of the data captured and the level of sharing required and is motivated by drawing analogies from Web analytics in which datasets are typically collected either server-side (e.g. through logging on the Web server) or user-side (e.g. by using modern analytics tools that execute within the user's browser such as Google Analytics) [KB00; Goo18c]. In digital signage, a similar categorisation can be applied: insights can be either collected by the infrastructure (server-side) or by the viewer (client-side), and the willingness to *share* collected datasets by the corresponding stakeholders (beyond their organisation) can be applied as an additional dimension. The intersections visualised in Figure 3.3 between the data origin and sharing dimensions are defined as follows:

Client-Side & No Sharing The analytics-relevant information is collected by the client (e.g. in the context of digital signage by the viewer) and is not shared with other stakeholders or organisations. Instead, datasets are made available to single stakeholders or organisations only. The browser cookie is the analogy from the web: it serves as a "small piece of data" [MDN18] created by a server and stored locally on the user's device and that can be retrieved by the creating server and domain only [MDN18].

Client-Side & Sharing The collected datasets are similar in their nature to the *Client-Side & No Sharing* category, however, clients are willing to share insights across multiple organisations. For example, two distinct stakeholders have access to information that was captured by the client such as display interactions. This intersection draws on analogies from advertisements in Web analytics in which certain Web sites might share page visit and interaction statistics with advertising networks to potentially improve advertisement sales, or simply as a proof that certain advertisements were shown.

Server-Side & No Sharing Insights are collected within a specific organisation or single stakeholder and are not shared with other entities. In the context of digital signage analytics, server-side sharing summarises insights that are collected by display owners, space owners and content providers and can include simple face counts, logs of content shown on displays, and retail-related analytics. The server-side category is analogous to Web logs stored on the Web server and capturing simple page visits.

Server-Side & Sharing The information collected server-side is shared across multiple distinct stakeholders and organisations. For example, space and display owners share their insights about viewer movement patterns with each other to improve the viewer experience by finding more appropriate locations for a display. The *Server-side & Sharing* category draws an analogy with industrial supply chain management and business to business relationships in which supply levels are shared to other organisations as described in more detail in Section 3.2.4 (Opportunities from Analytics Synthesis).

The matrix-based categorisation of collected datasets provides researchers with a framework for exploring the opportunities of data sharing across relevant stakeholders and organisations by specifically highlighting the intersections between the groups of stakeholders and the data origin.

3.2.4 Opportunities from Analytics Synthesis

Building on top of the data sharing matrix (Figure 3.3) and data collection overview (Table 3.1), we identified a set of example opportunities that arise from combining datasets across multiple distinct stakeholders. We additionally provide an overview of the feasibility of data sharing between individual organisations by presenting use cases from other domains in which data sharing is successfully practised.

3.2.4.1 Creation of Novel Insights

The creation of novel insights in the form of datasets and analytics reports can be seen as the main advantage of data synthesis across stakeholders. Examples from related domains such as supply chain management have shown that stakeholders are likely to share their analytical insights—under the condition that there is a clear benefit to all parties involved [LW00]. In the context of public display deployments, we have developed an initial set of examples of data synthesis across stakeholders based on the identified dataset individual stakeholders are capable of collecting (Section 3.2.2 – Stakeholder Data Collection Analysis). In the following subsections, we preset a set of opportunities for the creation and computation of reports by combining data initially collected and owned by two distinct stakeholder groups.

Display Owners and Space Owners Both stakeholders can complement each other's insights in the areas of gesture based user interactions and behaviour analysis of passers-by—yielding a comprehensive set of insights about the audience. This includes in particular questions on how displays affect the mobility patterns of viewers and their activity, and the proportion of people present in a space interacts with the display and its content. By additionally combining space-specific data including such as can be collected in a retail context (e.g. sales statistics), the stakeholders involved can start computing correlations of viewer activity and interactions with purchasing activities to, for example, compute a conversion rate.

Display Owners and Content Providers Content providers can contribute detailed insights about the content that is supplied to a display, and potential follow-up interactions of users with the distributed content, e.g. by subsequently visiting a Web site. Display owners have detailed anonymous counts and behavioural insights about the audience that potentially views a piece of content in the vicinity of a display. Combining datasets from both stakeholders enables content providers to retrieve insights about the audience and their demographics that has viewed a particular piece of content, while display owners gain details about the type of content shown and potential proportions of viewers following-up beyond the direct interaction with the display. In particular, display owners will be able to understand which kinds of content have lead to higher dwell times, and which demographics have responded to and engaged with a piece of content. These insights can contribute to the development of automated display schedules that dynamically adjust the content based on the audience—leading to an improvement both for the content providers by allowing them to improve their reach and for the display owners by increasing the value of their displays based on the currently present audience.

Display Owners and Viewers Combining datasets from both display owners and viewers enables us to answer questions regarding the number of returning viewers for individual display deployments (i.e. viewers provide us with the ability to identify returning visitors). In addition, insights directly collected from the viewer's perspective (e.g. through wearable cameras) can be combined with data from the infrastructure to capture objects and places viewers have seen immediately before and after walking by or interacting with a display.

Space Owners and Content Providers Content providers benefit from the space owners data in particular regarding the space owner's capabilities to count and track users beyond the vicinity of the display—allowing the content provider to understand, for example, how their content has influenced the viewer's behavioural patterns. In particular in the retail environment, the combination of content insights and sales statistics could lead to a new understanding of the effectiveness of advertising campaigns.

Space Owners and Viewers Space owners are limited in their ability to recognise and identify individual visitors and viewers present in their space whilst viewers complement the insights about their activities and behavioural patterns from their perspective – in particular beyond the boundaries of the space in which the space owner is able to capture relevant data. Insights through the combination of both datasets allow us to answer questions on the viewer's activities before entering the space, e.g. from which spatial and geographical location viewers originated.

Viewers and Content Providers In the collection of relevant analytics data, viewers and content providers complement each other in their capabilities to identify the viewer activities in the vicinity of the display and beyond displays. This provides insights for the content provider into the viewer activity patterns after interacting with their content, whilst viewers can benefit from more appropriate and targeted content shown on displays as they walk by based on their mobility and activity patterns.

The examples described above are based on the synthesis of datasets from two distinct stakeholder groups only. By increasing the numbers of stakeholders who contribute their datasets, the number of insights that can be gained grow substantially. For example, bringing together space owners, display owners and content providers would enable comprehensive insights into the influence of both displays and content to an audience across the entire space in which the display is deployed.

3.2.4.2 Additional Benefits

In addition to the creation of novel insights, further benefits can be gained from combining and synthesising analytics data across multiple distinct stakeholder groups. Phan, Kalasapur, and Kunjithapatham have shown the opportunities in a different domain: readings from multiple sensors can be fused with the goal of ensuring the accuracy of individual sensor readings and this can lead to more valuable results [PKK14]. Tian et al. show that by combining previously collected data on viewer behaviour (in this case, the content that the viewers have seen on a display) with additional contextual data (such as the display location and weather), the authors have shown an increased accuracy of their target advertisement system [Tia+12].

Additional benefits also include the inclusion of viewers as equal stakeholders of the overall analytics ecosystem – viewers can supply information instead being a subject observed by other stakeholder infrastructure. This approach also allows viewers to opt-out of data collection whilst data collection by other stakeholders through infrastructure sensors such as CCTV do not allow viewers to trivially opt out. The sharing of analytical insights, in particular if similar datasets are recorded, can also lead to cost savings—instead of two stakeholders capturing identical insights (such as audience insights through video analytics), the datasets can be captured once and shared across interested stakeholder groups. Alternatively, similar datasets collected multiple times can potentially increase the veracity of the insight by allowing stakeholders to cross validate the collected information.

Figure 3.4: Focus on the collection of traditional, *display-oriented sign analytics* data fundamental to the creation of viewer-centric analytics.

3.2.4.3 Attitudes Toward Data Sharing

Analytical insights such as sales statistics are often considered commercially sensitive, likely introducing significant burdens to sharing and combining such datasets.

While we are not aware of previous work that has explicitly explored the data sharing and synthesis across digital signage stakeholders, other examples in business-to-business relations exist. Examples include data sharing along supply chains where it is common to share information with business partners - providing insights and benefits to all those involved [LPW04; LW00]. Such data can specifically include inventory levels and product sales statistics automatically retrieved from point-of-sales terminals and transmitted to suppliers and other stakeholders in the supply chain [KLO04]. Kulp, Lee, and Ofek describe the specific benefits for manufacturers from entering such a data sharing agreement with retailers. Data from multiple stakeholders along the supply chain is combined to create better demand prediction for specific products (by analysing current and historical sales patterns), enabling manufacturers to dynamically adapt to current sales patterns and ultimately improve the performance [KLO04]. Both historical as well as real-time sales statistics, however, are commercially sensitive data and require written agreements and contracts between the involved parties prior to sharing as identified by Ghosh and Fedorowicz [GF08]. Kulp, Lee, and Ofek point out that at a most basic level, the requirement for entering data sharing and synthesis agreements is a clear benefit to all parties involved in the process [KLO04].

3.3 Capturing Traditional Signage Analytics Data

3.3.1 Overview

It is essential for signage analytics systems to capture data from a sign itself such as a log of the content shown and the physical state of the display in order to be able to combine these insights with additional metrics such as viewer mobility patterns as shown in Figure 3.4. In this section, we focus specifically on the collection of *display-oriented sign analytics* data. Current sign analytics systems (previously described in Section 2.3.1) are typically created specifically for the digital signage domain. This means, such systems are unable to leverage

Table 3.2: Example extensions to UMP to support *Proximity* and *Interaction* event types.

Hit Type: Proximity	Hit Type: Interaction
Event timestamp	Event timestamp
Referrer	Referrer
Content	Content
TID	TID
Queue time	Queue time
Session identifier	Session identifier
Client identifier	Client identifier
Total number of viewers	Interaction type
Number of viewers facing the display	X-coordinate of touch event
Number of viewers interacting	Y-coordinate of touch event
	Data precision (data collection frequency)

the wealth of analytics tools deployed for other domains such as the Web. To explore whether such existing analytics tools could be repurposed for the signage domain and to provide us with the ability to capture and process display-oriented analytics data, we designed and developed Pheme, an analytics support platform. Pheme consists of two core components: client-facing libraries for capturing relevant events on the display side, and a back-end system for processing, storing and aggregating incoming analytics data.

3.3.2 Client-side Data Collection

In order support the capture of display-oriented analytics data, we defined a data model and developed a set of client-libraries that provide the ability to describe and collect a standard set of traditional signage analytics metrics.

3.3.2.1 Data Modelling

To model signage analytics data, we draw on our literature survey (Section 2.3.1 (Systems) and existing analytics engines from the Web, utilising the Universal Measurement Protocol (UMP) – initially developed by Google for Web and mobile phone related analytics [Goo18h]. UMP defines attributes for a typical set of analytics events such as page views, e-Commerce related activities and the more general ‘events’ type that can be used to describe in-page or in-app viewer behaviour. In the context of digital signs, we describe the content shown on displays as a *Page View* hit type described as part of UMP, similar to a user opening a Web site on their browser. Typical attributes of the page view hit type include the description of the content shown, referrer attributes and details about the user agent (in this case, about the display that showed the content). To capture the physical state of the display, we utilise the custom *Events* type that provides a certain level of flexibility in describing analytics events using the fully customisable category, attribute, label and value fields. The physical power state of the display is categorised as `display_power_state` with the actual power states stored in the attribute field (e.g. ‘on’, ‘off’ or ‘unknown’). A more detailed mapping from signage-related analytics data to Web analytics is described in Chapter 4 (Reporting).

To illustrate the extensibility of UMP and to allow us to capture additional user-related interaction activities such as explicit interactions (e.g. through direct touch), and implicit interaction through display personalisation systems such as Tacita [Dav+12], we extended UMP with additional data attributes. Table 3.2 shows the set of attributes that are supported by Pheme to collect *proximity* and *interaction* related analytics events.

3.3.2.2 Display Client Libraries

We developed a set of front-end libraries that provide support for the integration of Pheme into existing systems and deployments. These libraries follow a similar principle that can be found in Web analytics systems which typically feature a set of simple libraries and code snippets for administrators to integrate on their Web pages to immediately enable basic analytics tracking [Goo18e]. These code snippets are automatically loaded as part of the interactive JavaScript code executed on the visitor's browser and report the visitor's activity automatically back to the analytics system. In digital signage, however, analytics need to be captured within the digital signage players—typically developed in proprietary software in the form of heavy clients. Due to the lack of standards and heterogeneous implementations, an automated or dynamic integration of an analytics client library to provide an analytics tracking automatised to the same extend as in the Web, we developed separate client libraries that provide developers with a way to easily integrate and enable analytics tracking within their infrastructure.

The Pheme client libraries are available as a set of Python-based modules (both for Python 2.7 and 3). Providing support for front-end systems and enabling the capture of user interactions, we also provide libraries for PHP and JavaScript (e.g. for dynamic, Web-based display content). The libraries provide methods that can be reused by developers and content creators to associate user interactions with appropriate UMP attributes. For example, direct user interactions (e.g. touch displays) can be captured through the integration of Pheme within a Web-based content item, while personalisation requests and contextual information can be captured and reported on the back-end.

An example code snippet for the integration of Pheme analytics tracking into a Python code base can be found in Listing 1. Developers are simply required to initialise `PhemeAnalytics` with their associated TID and can use the Pheme analytics instance to report events. Providing a unique user identifier through the `.set_client_uuid` method can be used to associate an event to a specific user, whilst reusing the same user identifier enables the tracking of users across across multiple instances and devices. Pheme libraries for Python 2.7, PHP and JavaScript provide an identical set of functionalities and tracking methods for developers.

If events are collected using Pheme's analytics libraries, each event is automatically associated with a timestamp at the point at which the event was reported, and with a timestamp at the point at which the event reached the Pheme back-end—in line with the industry standard of other analytics libraries [Goo18c]. In addition, display identifiers and TIDs are used to accurately determine the source of the event and enable us to filter out potential false reports.

```python
from phemelibrary import PhemeAnalytics

# Set up the tracking ID to be attached to all reports.
pheme = PhemeAnalytics("example_tracking_id")

# Set a client identifier for the display or user interacting.
pheme.set_client_uuid("client_uuid")

# Report a custom event asynchronously
pheme.track_event_async("category", "action", "value")

# Report a proximity event asynchronously.
pheme.track_d_proximity_async(
    distance='1.2m', person_uuid='xyz', time_spent='17s'
)
```

Listing 1: Pheme library example code snippet to enable analytics tracking for Python-based applications..

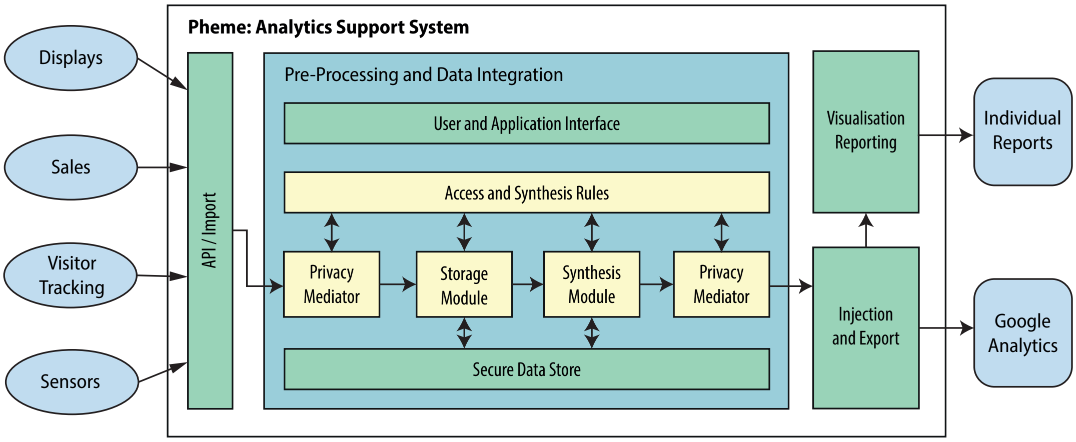

Figure 3.5: Pheme high level architecture and data flow diagram. Components highlighted in green have been fully implemented whilst components highlighted in yellow have not been implemented.

3.3.3 Server-side System Architecture

Once analytics data has been reported by the client-libraries using the appropriate function calls (Listing 1), it arrives at the Pheme back-end. The back-end system architecture has been designed following the pipeline concept and consists of four main components (Figure 3.5): *Import*, *Pre-processing and Data Integration*, *Export*, and *Visualisation and Reporting*.

3.3.3.1 Data Capture and Import

The *Import* module provides a RESTful API to which analytics-relevant events are reported from client devices using the UMP format. In line with common analytics systems, incoming events must consist of a unique tracking identifier linking incoming requests to a specific owner. The tracking identifier can be retrieved through a Web-based interface after registering with Pheme (Figure 3.6a).

3.3.3.2 Pre-processing and Data Integration

Following the import, the incoming datasets are passed on to the *Pre-processing and Data Integration* module. Pheme has been designed to support multi-stakeholder data synthesis as initially described in Section 3.2.4 (Opportunities from Analytics Synthesis) and to accommodate individual stakeholder requirements regarding privacy and data ownership [MCD17]. Both the data synthesis and the accommodation of privacy and data ownership requirements is performed within the pre-processing and data integration module.

This module consists of the following components: *user and application interface*, *access and synthesis rules*, *secure data store* and a set of processing components (Figure 3.5). Incoming datasets from the Import module arrive initially at the in-bound *Privacy Mediator* component (based on initial ideas first proposed by Davies et al. [Dav+16]). This module provides stakeholders with a mechanism to filter sensitive information or anony mise data prior to sharing and synthesising with datasets owned by other stakeholders – e.g. to protect business secrets whilst still contributing to the generation of novel analytical insights if data has not been anonymised prior to sending to Pheme. Following mediation, incoming data is then processed with respect to predefined processing and synthesis rules supporting a fine-grained level of control regarding the processing and sharing of data. The *Synthesis Module* is the core component for the synthesis and combination of mediated datasets in accordance to the rules specified by relevant stakeholders. Upon the successful processing and creation of synthesised and combined datasets, both mediated and synthesised datasets are stored in the *secure data store* for potential future use, e.g. further data synthesis for other stakeholders and allowing future reference to historical datasets and supporting a full range of (future) synthesis and redaction policies.

To make the synthesised datasets available to other stakeholders, the architecture consists of an outbound *Privacy Mediator* module that applies policies specified by contributing stakeholders. For example, some stakeholders would agree to provide more insights to enable a better data synthesis and linkage with other stakeholders' datasets but under the condition that certain association parameters (i.e. parameters used to combine multiple distinct analytics datasets such as the physical location of a display relative to purchasing statistics) used will be removed prior to releasing synthesised insights. Upon passing through the outbound privacy mediator, data streams are passed on to the injection and export modules for further processing.

3.3.3.3 Injection Modules and Export

To support the ability to repurpose existing third-party analytics engines, we designed the *Export* module that consists of a set of *injection components* that map and convert processed data to become compatible with the target system and 'inject' the converted data objects into third-party systems using the target systems' APIs. The Export module supports multiple "injectors" – each representing a specific mapping of data to make it compatible with the third-party service and includes an implementation of third-party APIs to be able to report relevant datasets. Users of Pheme can assign multiple injectors to their unique tracking

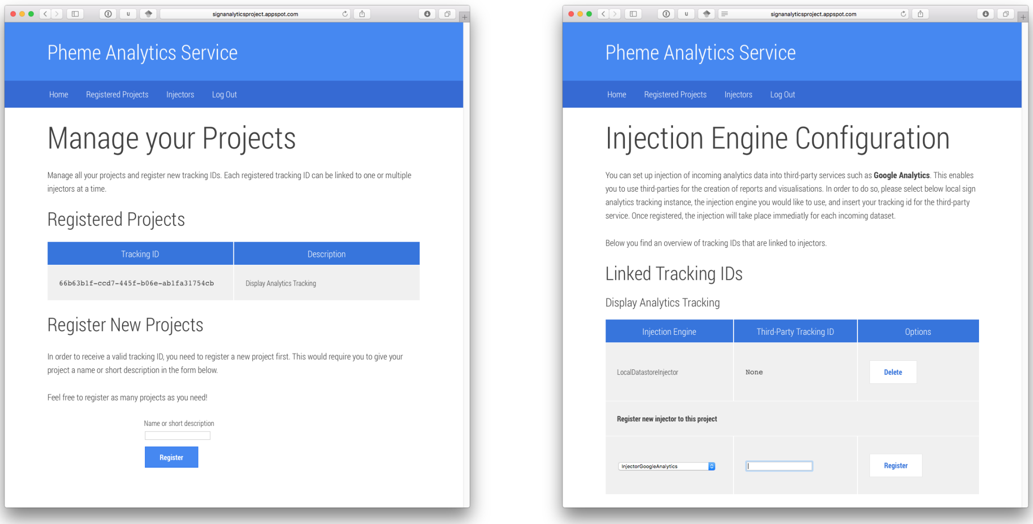

(a) Pheme User Interface: Overview of registered projects.

(b) Pheme User Interface: Overview of available injector modules.

Figure 3.6: Pheme user interface.

identifiers, therefore being able to automatically distribute incoming analytics data using different injectors simultaneously to the same analytics engine using different mappings, and to distinct analytics services.

As third-party analytics systems are typically developed for a specific purpose (e.g. for Web analytics) and only support the reporting and visualisation of specific metrics, Pheme supports the creation of custom reports and visualisations that go beyond the capabilities of individual third-party analytics engines. For example, users can choose to implement custom reports and visualisations within Pheme with direct access to datasets associated with the users' tracking identifier in addition to injecting their datasets into third-party analytics systems.

3.3.4 Implementation

Significant portions of Pheme including client libraries, a user interface for administrative purposes and the back-end system have been implemented – an overview of the implemented components is provided in Figure 3.5.

3.3.4.1 Client Libraries

The described client libraries have been implemented both for Python 2.7 and Python 3 and have been integrated in the e-Campus display testbed. Additionally, we developed client libraries in PHP and JavaScript for their use as part of interactive, Web-based display applications to capture explicit or implicit user interactions and system events.

3.3.4.2 User Interface

We implemented a set of Web-based user interfaces through which users can register to collect analytics data by creating a new unique user identifier and assigning the required set of injection modules—similar to configuration interfaces that can be found in other (Web) analytics systems. Figure 3.6 shows an example of the user interface that users can access (Figure 3.6a for the registration of new tracking identifiers; FIgure 3.6b for the configuration of injectors).

3.3.4.3 Server-side System

We implemented core components of the Pheme back-end architecture including the Import, Injector and Export components as well as a basic Pre-processing and Data Integration module. Pheme has been initially implemented entirely in Python on top of Google's App Engine framework and runs in its entirety on Google Cloud services. The utilisation of existing processing power in the cloud enables us to dynamically adapt to high server loads and cope with potentially large incoming datasets and a high number of parallel requests. In particular, processing components of Pheme (i.e. import, parsing of incoming requests, and the application of appropriate mapping and export) are placed in separate process instances within Google App Engine, and the number of instances automatically increases with the number of incoming requests and the associated process load. The storage components are separate to the processing components – we utilise Google Cloud Datastore that, similarly to App Engine process instances, automatically scales and adapts to the incoming load of read and write requests.

3.3.4.4 Evaluation

We note that the evaluation of Pheme can be found in Chapter 6, Section 6.2.

3.4 Capturing Viewer Mobility Data

In the previous section, we considered capturing display-oriented data directly from the signs. In this section, we focus on the collection of viewer-oriented mobility data to complement the remainder of the required set of analytics data (Figure 3.1). To explore the collection of viewer mobility data, we have considered three distinct mechanisms:

1. **viewer-based tracking** using a mobile phone application (Section 3.4.1, Viewer-based Tracking),
2. **infrastructure-based tracking** of viewers using Wi-Fi fingerprinting (Section 3.4.2, Infrastructure-based Tracking), and
3. **synthetic mobility traces** using mobility models to generate viewer-oriented mobility data (Section 3.4.3, Synthetic Analytics).

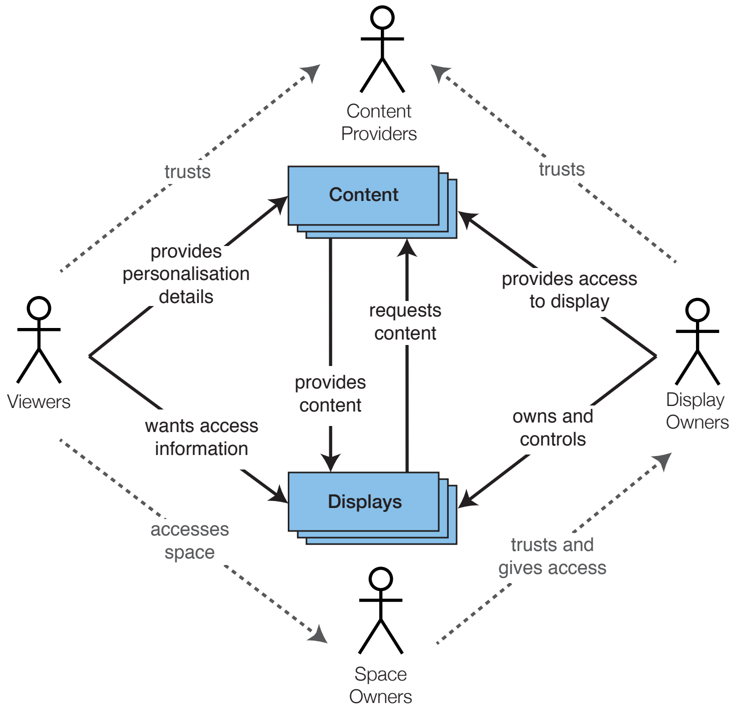

Figure 3.7: Tacita trust relationship between stakeholders of open pervasive display networks – adapted from Davies et al. [Dav+14].

Both the viewer- and infrastructure-based tracking approaches provide us with the ability to collect detailed viewer mobility traces and understand their interactions and engagements with signs beyond the scope of single display deployment. Due to the high sensitivity of the involved datasets, we developed the *synthetic analytics* approach as a third alternative to generate and collect comprehensive viewer-related mobility traces based on predefined models.

3.4.1 Viewer-based Tracking

Our work on viewer-based tracking builds on top of the Tacita display personalisation system proposed by Davies et al. [Dav+14]. Whilst in typical display personalisation systems, providers and large commercial signage networks utilise video analytics and other techniques to track individuals and tailor the content towards them, the approach initially suggested in [Dav+14] builds on top of trust relationships of viewers with personalisable services and content providers – by explicitly disconnecting the viewer from infrastructure providers such as organisations owning large digital signage networks [Dav+14]. Figure 3.7 visualises the core concept of the architecture in which display owners, content providers and users are three separate entities. To respect viewer expectations of privacy, both viewer locations (to determine their proximity to displays) and their preferences (to determine the kinds of content they wish to see) are only shared with a trusted application the user prefers to personalise – and are not shared to the signage network. As a result, the architecture protects sensitive data and prevents signage network providers collecting large quantities of location traces and

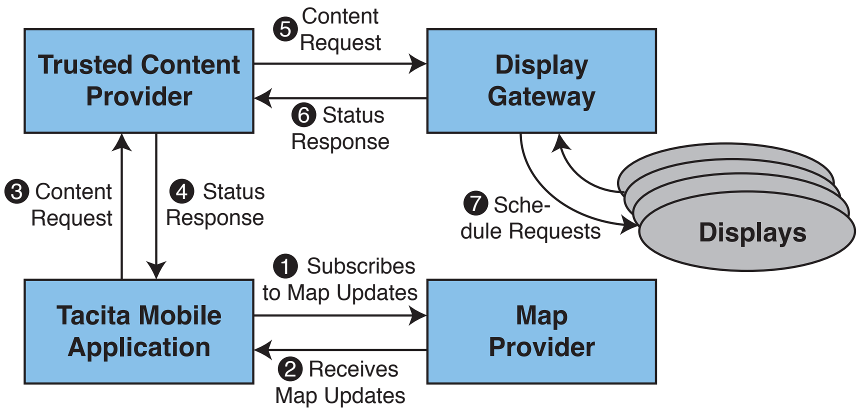

Figure 3.8: Tacita system architecture and data flow diagram (based on [Dav+14] and adapted in [Mik+18d]).

personalisation preferences whilst still improving the user experience by supporting display personalisation.

For our work on digital signage analytics, Tacita provides us with a way to collect anonymised display sighting events and personalisation preferences. Detailed analytical insights that have been captured and used in the context of Tacita and as part of a self-learning digital signage platform are described in Chapter 4 (Reporting), and the use of Tacita for actuation purposes are described in Chapter 5 (Automated Use of Pervasive Display Analytics) while in the remainder of this section we focus on the use of Tacita for data collection.

3.4.1.1 Back-end Systems Architecture Design

We have restructured and redesigned the systems architecture initially proposed by Davies et al. [Dav+14] and adapted it to the needs and requirements of in-the-wild displays deployment and state-of-the-art location tracking systems. The Tacita architecture developed in the context of this dissertation consists of five core components (Figure 3.8):

Trusted Content Provider Personalisable applications (described as *Trusted Content Providers*) are trusted both by the viewer and the display infrastructure and are designed to provide personalisable content for viewers through public displays. Viewers can access trusted content providers through a dedicated application on their mobile phones and provide their preferences. In addition, trusted content providers also serve the content that is shown on displays. Typically, trusted content providers serve dynamically generated Web-based content depending on the contextual information and preferences of the requesting passers-by.

Display Gateway The *Display Gateway* serves as a gateway for connections to individual display nodes that belong to a single organisation and might be located behind a firewall or inside a closed network. The gateway protects individual display nodes from being exposed on to a wider network and provides a public facing application programming interface as typically individual display nodes are secured behind firewalls and within private networks to minimise the attack surface. To allow content personalisation, however, display nodes are required to provide appropriate interfaces for third-party applications. Requests for dynamic changes of content on request (typically based on detected viewer presence) are sent by the trusted content provider to the appropriate Display Gateway depending on the viewer's location instead of communicating directly with the display node that is closest to the viewer. trusted content providers request screen time by providing the location, display and application identifiers to the gateway which, in the background, maps the provided identifiers onto a physical display and negotiates the (immediate) transition of content with the individual display node. Whilst the interface of the display gateway facing the trusted content provider is standardised, the communication with individual display nodes can be proprietary and is likely to be specific to individual public display networks.

Public Displays Individual display nodes serve as 'regular' public displays and, depending on their configuration, cycle through their regular display content schedule. Further, display nodes support an interface to retrieve messages from Display Gateways to change their current display schedule on demand. After receiving a content scheduling request from the Display Gateway, display nodes access the content from the Trusted Content Provider—typically in form of a dynamic, Web-based application. The actual content delivery with respect to individual viewers present in the space is conducted by the trusted content provider exclusively. To enable trusted content providers to map the requesting display to viewer sightings of displays to deliver the appropriate piece of content, each display appends their unique location and hardware identifier to the `HTTP GET` request when opening the display client.

Map Provider Map providers describe distinct services that supply maps and descriptions of displays and personalisable applications. Maps are downloaded by the client applications installed on the viewers' mobile devices and updated on a regular basis.

Tacita Mobile Application Viewers can express their preferences for display applications and content through a dedicated mobile client. The client also allows viewers to explicitly configure their levels of (location) data sharing and the content provider that is trusted by the viewer.

This approach differs from the initial architecture proposed by Davies et al. [Dav+14] in which third-party personalisable applications directly communicate with individual display nodes to request a change of content and screen time. In the context of the existing networks of public displays we found this approach infeasible for a number of reasons. Firstly, individual display nodes are often situated behind a network firewall preventing any access to the screens from the outside world due to increased security purposes. Secondly, the proposed approach would require the introduction of appropriate interfaces and security measures on every single display node, limiting the scalability of the system. As a result, we made significant modifications to the original architecture and introduced the *Display Gateway* as a new component, shown in Figure 3.8. The Display Gateway provides a standardised single interface for third-party applications that can be accessed to request screen time and maintains a direct connection to display nodes part of the network. Incoming requests from a third-party application are first validated and verified, and then mapped and forwarded onto the corresponding display node. Each display network that becomes part of the Tacita framework needs to supply at least one Display Gateway that maintains a list of associated displays and supports the mapping of incoming requests to an actual display.

3.4.1.2 Mapping of the Public Display Network

In the Tacita architecture proposed by Davies et al. [Dav+14], displays announce their personalisation capabilities including personalisable and interface endpoints through appropriate protocols such as Bluetooth. The use of central repositories with maps of public displays was described as an optimisation removing the need for the viewer to be in proximity of a personalisable display in order to retrieve a list of applications and begin the configuration of preferences [Dav+14]. However, state-of-the-art BLE beacon protocols including iBeacon are not well suited to transmit large quantities of data such as application descriptions [Mal+16] and hence in our re-design we focused exclusively on maps.

To support this, we developed a unified JSON-based map schema to accommodate the description of personalisable display networks (i.e. multiple Display Gateways) and the collection of supported personalisable applications for displays part of the network. In addition, the map includes basic information about the set of displays with associated location metadata. To detect viewers in proximity of displays, Tacita requires for each display the specification of *Trigger Zones*, i.e. spatial areas in which the entrance and exit of viewers is detected and the associated display is requested to dynamically change the content. Conceptionally, trigger zones can be specified independent of potential technology choices using any proximity- or geofence-based descriptions. As part of Tacita, location information are typically expressed through proximity iBeacons[2] and absolute geographical locations in the form of latitude and longitude attributes. An example map description including a single display and personalisable display application is shown in Listing 2.

3.4.1.3 Mobile Applications

The Tacita mobile application is the only component visible to the end user and serves as a platform to express personalisation preferences (Figure 3.8). The mobile phone client consists of two main features:

1. the detection of nearby displays in accordance the specifications supplied by the map provider (i.e. in this case, the specification of BLE identifiers), and

[2]iBeacon is a Bluetooth Low Energy protocol developed by Apple and mainly used for indoor localisation.

```
{ "id": "24061166-c3ea-45e3-afc7-c027f9e82fdd",
  "meta": {
  "description": "e-Campus Displays",
    "start_date": "02/12/2017",
    "expiration_date": "02/01/2018",
    "publication_date": "02/12/2017",
      "map_version": "1.2"
    },
  "domains": [{
    "name": "Lancaster University Campus",
    "proxy": "https://example.com/content_request",
    "entries": [{
      "user_region_as_triggerzone": false,
      "regions": {
        "circular_regions": [{
          "lat": "54.01093",
          "long": "-2.78445",
          "radius": "30m"
          }]
      },
      "trigger_zones": {
        "circular_regions": [],
        "proximity_beacons": [{
          "beacon_major": "24",
          "beacon_minor": "6",
          "beacon_type": "iBeacon",
          "beacon_uuid": "d8e40b29-7649-428e-b80c-ba3ed0911fb4"
        }]
      },
      "capabilities": {
        "uuid": {
          "display_id": "display-7",
          "display_name": "Display Foyer 7"
        },
        "display_apps": [{
          "name": "Bus Departures",
          "callback_url": "https://example.com/tacita_callback",
          "description": "Bus Departures description",
          "icon_url": "https://example.com/tsp_bus_logo.png",
          "homepage": "https://example.com",
          "config_url": "https://example.com/config"
        }]
      }
    }]
  }]
}
```


Listing 2: JSON map description (initially published in [Mik+18d]).

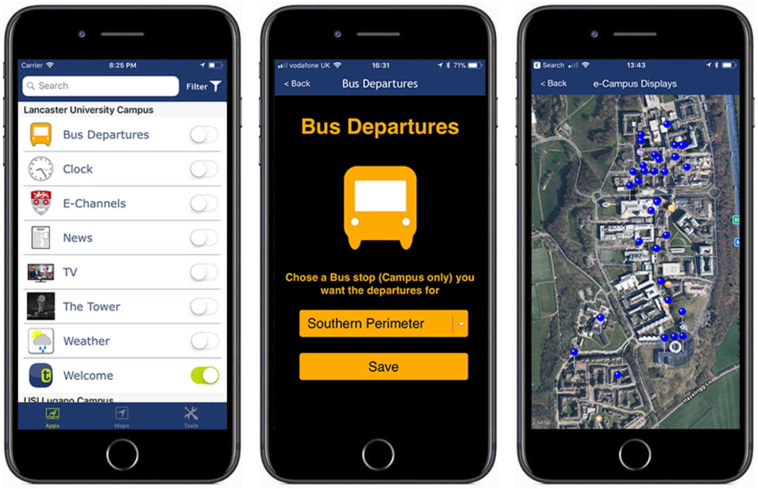

Figure 3.9: Tacita mobile application for Apple iOS (initially published in [Mik+18d]).

2. providing viewers with a utility to specify and express their preferences and initialise trust relationships among a set of personalisable applications (Trusted Content Providers).

Figure 3.9 provides an overview of the basic features provided by the mobile phone client to the user. After the first use of the application, viewers are presented with an overview of all available Trusted Content Providers for personalisation (Figure 3.9, left). Each Trusted Content Provider can be activated and deactivated at any time. Activating a Trusted Content Provider leads to the initialisation of a trust relationship between the user and the selected Trusted Content Provider – as from the point of switching on the viewer's proximity to displays will be shared with that particular Trusted Content Provider. The use of an integrated Web-based configuration screen enables Trusted Content Providers to serve specific pages for their application, hosted on their premises and even enable the Trusted Content Provider with the ability to support authentication mechanisms such as OAuth2.

3.4.1.4 Opportunities for Data Collection

Tacita provides us with a unique platform for the collection of analytics regarding viewer interactions and navigation patterns. In particular:

Trusted Content Provider Configuration As configuration pages are served from individual Trusted Content Providers as Web pages, we collect user interaction on the Trusted Content Provider through simple Web analytics and server-side access logs. This includes the preferences specified by the user and frequencies of configuration page visits.

Display Proximity We are able to capture viewer 'to' display proximity events each time the viewer enters a trigger zone. These logs include both data captured on the viewer's mobile phone client, and server side. First, the Tacita mobile phone client captures the timestamp of a viewer entering (and exiting) a trigger zone and supplements the information with metadata associated with the trigger zone. After sending the display proximity event to the Trusted Content Provider (in the form of a content request), the Trusted Content Provider back-end captures and stores the incoming request, and supplements the information with the timestamp of the incoming request. This enables the Trusted Content Provider to capture potential latencies and delays (under the assumption the client device is time synchronised). In accordance with the system architecture (Figure 3.8), each incoming enter-report is immediately forwarded from the Trusted Content Provider to the display gateway whilst removing any user identifiable information.

Display Content As part of Pheme (Section 3.3: Capturing Traditional Signage Analytics Data), we capture the content displayed of any display part of the signage network — including Tacita requests and are therefore able to tell which content the display was showing at the time the viewer was detected in proximity of a display.

All datasets including user identifiers are collected on the individual Trusted Content Providers that viewers have activated. In addition, we log the behaviour and response rate of display nodes to personalisation requests. This includes each incoming request on the display gateway, both successful and failed messages to display nodes. In addition, Trusted Content Providers can track whether their content was successfully shown by monitoring access to the display client pages that are used to deliver content to displays.

3.4.1.5 Implementation

The back-end components of Tacita including the Display Gateway and an initial set of Trusted Content Providers (such as personalised weather forecasts and travel) have been implemented entirely in Python. The Tacita mobile phone client has been developed in Swift and made available for iOS-based mobile devices in the Apple App Store[3]. We have integrated Tacita into the e-Campus display deployment and made it available on over 45 displays across the University campus. We have further expanded the deployment to the campus at the University of Lugano in Switzerland.

To support the proximity detection of viewers near displays, we have utilised Bluetooth Low Energy and Apple's iBeacon standard. Whilst in principle trigger zones can be set to arbitrary sizes, in practise the upper bound of the trigger zone size is mandated by the range of the BLE-powered iBeacons and viewers' mobile devices (typically around 10 meters) or the type of (indoor) location tracking that has been used.

3.4.1.6 Evaluation

We note that the evaluation of this system is described in Chapter 6 (Trials), Section 6.3.

[3]Tacita can be downloaded at https://appstore.com/tacita.

3.4.2 Infrastructure-based Tracking

While client-side tracking enables the creation of rich analytics, it is also possible to capture analytics data using an infrastructure-based approach. Similar to client-side tracking, the overall aim of infrastructure-based tracking in the context of digital signage is to capture contextual information about viewer navigation patterns and to support display personalisation. This is an alternative approach to collecting viewer mobility data as we proposed as part of Figure 3.1. In this section we describe the most commonly used approach of infrastructure-based tracking in the form of Wi-Fi fingerprinting – a commonly used technique across commercial environments such as shopping malls. The advantage of infrastructure-based tracking is that viewers are not required to install a dedicated application on their mobile phones, and the environment is not required to be equipped with additional tracking technology. Typically, data collected in the environment is used to understand common walk paths and navigation patters inside the enclosed environment [Kha+13; Jay+16].

We have designed an infrastructure-based display personalisation system in the context of LiveLabs, "a first-of-its-kind testbed that is deployed across a university campus, convention centre, and resort island and collects real-time attributes such as location [...] from hundreds of opt-in participants" [Jay+16]. In particular, LiveLabs has been deployed at a large convention centre located in Singapore, collecting location information of all visitors present in the space with a frequency of approximately 15 seconds. LiveLabs consists of an event-based API that allows external services to access current location information at the point at which they have been captured.

3.4.2.1 System Design and Architecture

The system architecture consists of six modular components that have been based on Tacita (Section 3.4.1): *context data fetcher*, *pattern recognition*, *content creation and selection*, *infrastructure connectors*, *context store*, and a *real-time and event-based content trigger* engine.

Similar to Tacita, the system relies on proximity information of viewers to public display and the specification of trigger zones around individual displays which are used to detect viewer enter and exit events to change the content shown accordingly. Due to the availability of different mechanisms that can be used to detect viewer proximity, we included the *Context Data Fetcher* component that can be used to plug in any external source for location tracking including Wi-Fi fingerprinting provided by LiveLabs [Jay+16]. Each external data source is modelled as an individual process with access to a local context store through which location update events are communicated to other system components in real time. Further, external data sources are required to provide a unique user identifier and the users' location points as a basic set of information. The proximity detection of viewers to displays is computed within the *real-time content trigger* module: each location point of a potential viewer is compared with the set of trigger zones of displays to detect whether individual viewers are within a display trigger zone. The system utilises an event-based approach: real-time

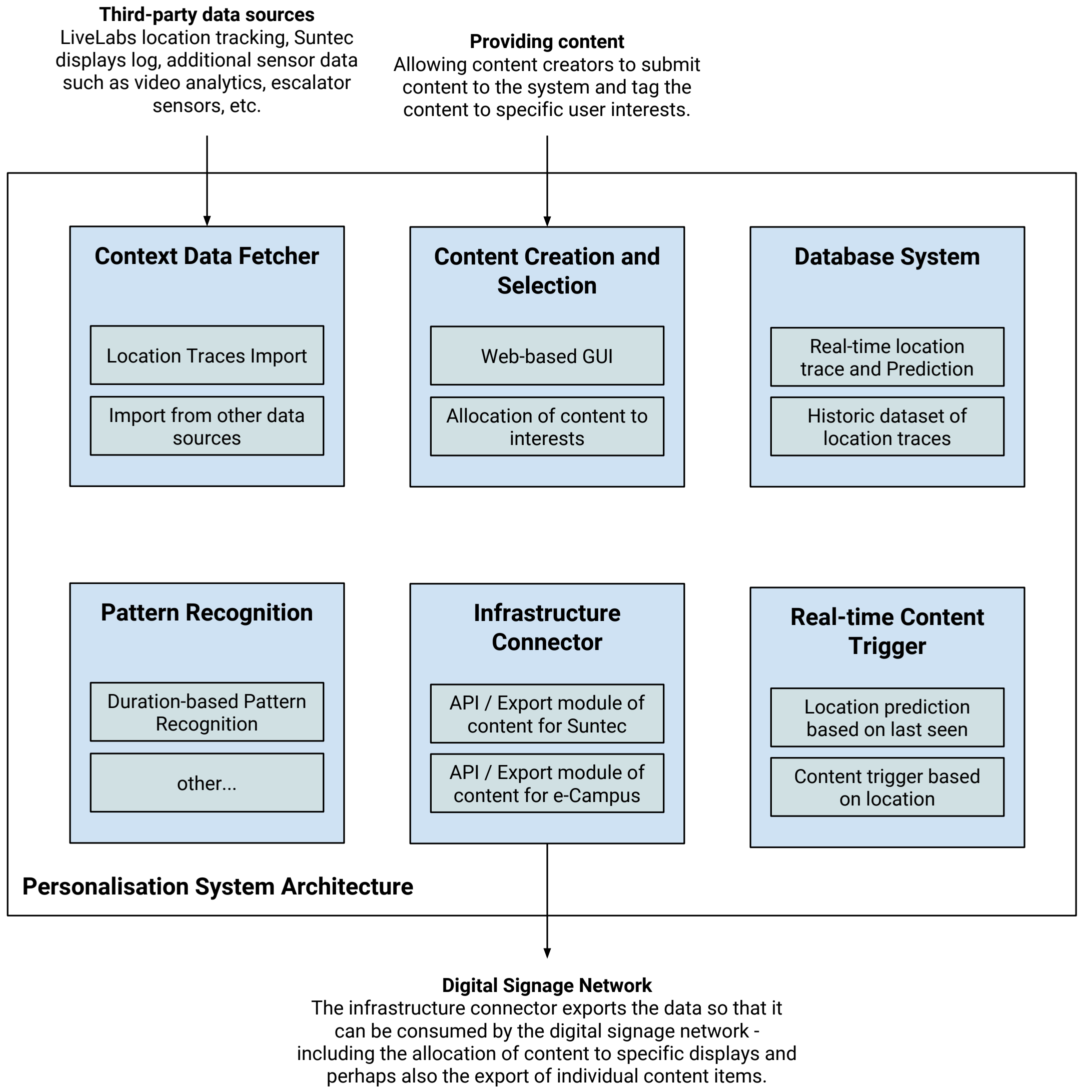

Figure 3.10: System architecture for infrastructure-based tracking and public display personalisation integrating LiveLabs [Jay+16] real-time location tracking.

location updates are passed on to the real-time content trigger component that detects whether viewers are in proximity to displays and finds corresponding rules for content selection for individual viewers. Content preferences are computed on a constant basis within the *pattern recognition* component that, based on historic navigation patterns stored in the data store, determines potential viewer interests—for example, based on duration spent in certain rooms and locations. The actual content and content selection rules are manually supplied by the space owner and display providers and defined within the *content creation and selection* module. Rules, for example, can consist of location and time dimensions. The *infrastructure connector* component is, similar to the *display gateway* in Tacita, responsible for the mapping of visitor proximity to the actual display and immediate content change requests on the display.

More details on the selection of content based on historical traces of viewers is described in Chapter 5 (Automated Use of Pervasive Display Analytics).

3.4.2.2 Opportunities for Data Collection

In the context of infrastructure-based viewer tracking, collected datasets exclusively originate from the space owner and can be categorised as *server-side no sharing* in accordance with our data categorisation framework described in Section 3.3. Whilst advantages of this approach include the absence of required viewer opt-ins and dedicated mobile phone applications, the quality and accuracy of using infrastructure-based sensing may appear to be of lower quality depending on the technology used.

However, Wi-Fi fingerprinting enables the collection of comprehensive navigation patterns of viewers within a space beyond the vicinity of the display and provide us with “a broader view of the surroundings” [Section 3.2.2, Space Owners] beyond the immediate vicinity of displays. More specifically, the collected datasets include anonymised user identifiers, location points within the space, date- and timestamps, and an accuracy measure. The rich dataset allows us to compute viewer proximity to displays, dwell times at any location within the space, and map user locations to specific conference venues and events. The resulting dataset gives us insights into potential interests of individual viewers: for example, a long dwell time in a certain conference room in the convention centre can be mapped onto a specific event and used later-on for personalised content selection. In addition to viewer-related data, we also capture content displayed on each digital sign within the venue that has been integrated in to our system. The resulting dataset of viewer mobility paths can be enriched by accessing Pheme to retrieve content shown and timestamps of associated displays.

3.4.2.3 Considerations

Whilst the collection of real-time location data is fundamental to this system, the use of the collected information for the benefit of the visitor is important regarding the acceptance of the system—in contrast to systems that collect mobility data or audience demographics purely for the analytics and the benefit of the space or display owner.

Using location traces to improve viewer experience requires the addition of a dimension that allows the association of content with displays *and* date and time attributes to show content in the future according to viewer presence in the past. To better understand which display content works best for viewers and improve content selection algorithms, the lack of a feedback loop (i.e. capturing viewer responses to content) provides an additional challenge in the design and development of potential machine learning algorithms and understanding which content worked best and was most effective. However, real-time location traces can be a suitable mechanism to identify groups of users to target and, in particular, automatically generate targeted content based on current contextual information from the vicinity of the display without using potentially more privacy invasive technologies such as face recognition or face identification.

3.4.2.4 Implementation

The system was integrated with LiveLabs [Jay+16] and the pervasive display deployment at Suntec Singapore [LG 13] consisting of a large high-definition video wall and over 100 digital signs located across the conference venue. The back-end components have been implemented entirely in Python whilst the front-end components and content player were written using Web technologies and AngularJS. The display network at Suntec is using the Four Winds digital signage platform[4]. Due to the limited functionality in the version deployed at Suntec, we implemented a dedicated Web-based public display client for dynamically displaying content that has been shown in full screen through the Four Winds display player. The client page and the back-end communicate through Web sockets, supporting real-time communication to immediately change the content based on viewer proximity.

3.4.2.5 Evaluation

We note that the evaluation of this system is described in Chapter 6 (Trials), Section 6.4.

3.4.3 Synthetic Analytics

In Sections 3.4.1 and 3.4.2 we focused on the collection of mobility data from using viewer-centric and infrastructure-based approaches respectively. However, the tracking of viewer mobility data is not always possible in certain spaces, e.g. due to viewers unwillingness of sharing relevant data, or the space not equipped with appropriate indoor location tracking systems. For example, Wi-Fi fingerprinting requires appropriate Wi-Fi base stations deployed throughout the space and time consuming configuration and initialisation. Additionally, the tracking of viewer mobility data and behavioural patterns can be considered as highly sensitive and privacy-invasive.

Following the viewer-centric analytics approach, we focus in this section on an alternative way to collect viewer mobility data in a privacy-preserving way that can be applied in any

[4] https://www.fourwindsinteractive.com

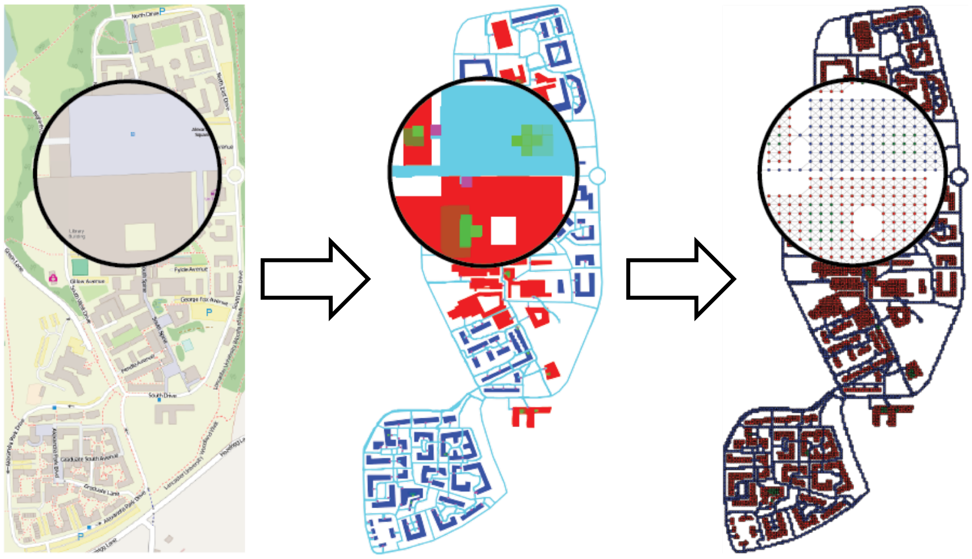

Figure 3.11: The basis for synthetic analytics are regular maps (left), which are annotated in a subsequent process (middle) and automatically transformed into a graph-based structure (right).

space and environment. Leveraging previous work in mobility simulation, we have designed a system that brings both mobility simulation and signage analytics together. The system is composed of three main components: 1. an underlying spatial map data structure with an encoded display view probability model, 2. a synthetic viewer mobility engine, and 3. an interface for the integration of real-world datasets. Figure 3.1 visualises the general core concept of the proposed approach: display-oriented analytics (e.g. content play logs) are combined with synthetic viewer mobility traces. The outcome are viewer-centric analytical insights on a comparable level as captured through viewer location tracking. We have designed a prototypical synthetic analytics system that supports the creation of mobility models in the context of the Lancaster University campus, and the combination with data collected as part of the e-Campus display testbed.

3.4.3.1 Spatial Map and Probability Model

The underlying spatial map forms the foundation for our synthetic analytics system. Maps are required to encode the spatial layout in which agents (i.e. instances of modelled viewers) will be operating and consists of the locations of displays, properties of viewer movements, and the probabilities of seeing a display and the content when walking by from individual spatial locations on the map (e.g. based on the 'viewable' areas of a display) considering the likely sight lines. We use a process in which we transform geographical maps into a graph-based data structure and encode additional attributes. Figure 3.11 visualises the general process in which the spatial map is encoded and translated into a undirected graph data structure on which the mobility model can operate. As the mobility of viewers and, in particular, display views are highly dependent on the locations of the displays, pathways, buildings and their entrance and exit points, and other properties, we have first encoded display locations, pathways and buildings using a simple colour scheme (Figure 3.11, middle). Each display location is represented by a green square whilst the probabilities of a passer-by present in the vicinity of the display for seeing and engaging with the display is encoded using the alpha value (e.g. the probability of 0.5% seeing a display from a node in a location corresponds to the alpha colour value of 0.05). This probability value builds on our experience with the digital signage deployment and has been chosen based on the placement and visibility of the individual displays. For example, a more prominent placement of a display (e.g. a display inside a building foyer facing an entrance) is more likely to be seen compared to other displays across the University campus. Other colours are used to encode buildings, pathways, entrances and exits. We use a simple approach that translates the colour-coded map into a corresponding set of graph nodes and edges in which each node corresponds to a spatial location of a building, pathway, entrance and exit, and display. Each nodes further include metadata such as the display identifiers and probabilities of seeing an adjacent display. Nodes are connected through edges on which the mobility model can operate and move.

3.4.3.2 Mobility Simulation

The mobility simulation component was designed to support an agent-based simulator in which each virtual person and viewer are individual instances of a model – similar to related work conducted in the area of mobility simulations [Bon02]. For the initial system, we have developed the following three simple mobility models to reflect the majority of the population typically present on the University campus.

On-Campus Student A model representing a student who lives on the premises of the university and randomly changes their location between colleges, departments, and lecture theatres once an hour during university core hours between 8.00am and 10.00pm. More specifically, the model finds a destination every hour at random and chooses a simple shortest path algorithm to move along the graph structure to the destination. The student starts and ends the day at the same college building that it was randomly assigned at the initialisation phase of the simulation.

Off-Campus Student A model representing a student who lives off-campus and commutes onto the university premises by bus – reflecting the mobility patterns of the majority of students. In this model, we assume that students arrive at a bus station that is located in the centre of campus at an arrival time chosen at random between 8.00am and 11.00am, and subsequently follow the identical movement patterns and strategies of the On-Campus Student model. The Off-Campus Student model chooses a 'leave time' between 3.00pm and 7.00pm at which time agents navigate back to the central bus station.

Random Building Navigator A baseline model in which potential viewers constantly move between two points on the graph chosen at random throughout the entire lifetime of the simulation. This model is an important comparison against any other model in the simulator as it is able to explore a wide spectrum of mobility scenarios without any limitations and assumptions of the other models regarding their movement patterns. The single criteria of this model are the movement times: the Random Building Navigator only moves between 8.00am and 10.00pm – following similar movement strategies as the other models.

The synthetic analytics system is designed to generate synthetic mobility traces by creating a pre-defined number of agents for each model (each instance can be only based on one model but the system supports a mix of multiple instances based on various models). To generate a realistic set of mobility traces associated to each agent, the mobility simulator performs a number of predefined iterations and allows each agent to move along the graph structure whilst agents are allowed to move to one neighbouring node that is connected through an edge in each iteration. Of course, due to the mapping of nodes to the spatial layout of campus, each iteration and movement of an agent from one node to another directly correlates to a virtual person moving across campus within the given time frame. Therefore, each iteration of the simulation translates to a time frame based on the preset parameters and size of the spatial base layout that has been translated into the graph structure providing us with the ability to associate a timestamp to each visit of an agent to a node if the start date and time of the simulation is known. The use of appropriate parameters is particularly important due to the graph structure and node directly correlating to the spatial map layout, i.e. edges between adjacent nodes directly translating to a distance in the underlying map.

In each iteration of the simulation, individual agents are given the option to change their position. For any agent, a range of tests are performed. If an agent visited a node within the viewable area of a display, the encoded probability value is used in order to determine whether the agent has glanced at the display and seen the content. Display glances are registered and tracked as part of the individual agent metadata and can be exported later-on separate to the entire set of visited nodes. The results are a comprehensive set of display glances and movement traces. Both the display sightings and the history of nodes that have been visited including associated timestamps are stored as part of the metadata of the agent. In addition, each agent includes of an universally unique identifier.

In contrast to the display sightings collected as previously introduced in Section 3.4 (Capturing Viewer Mobility Data), we additionally gain comprehensive movement patterns beyond the vicinity of displays for the entire (modelled) population of the University campus—without the invasion of their privacy.

3.4.3.3 Modelling Synthetic Viewer Mobility Traces

We note that to support the creation of mobility models for digital signage analytics, we draw on a body of work in the area of mobility modelling that has been a widely explored topic in the context of other research domains. Researchers have used people mobility models in the context of theme parks [Che+13; AML13]. Cheng et al. [Che+13] developed an "agent-based simulation approach" [Che+13] which primary use was to understand the flow of people in a theme park and find opportunities for improving visitor experiences, reducing potential bottlenecks in the flow of people across the park, and evacuation management [Che+13]. The authors have modelled the park as an undirected graph data structure in which each node illustrates an attraction, and edges are paths between attractions. Each synthetic person that moves across the park on the graph data structure is modelled as an individual and independent agent making decisions which routes to take and when to move based on an underlying movement model [Che+13]. In a similar context, Aravamudhan, Misra, and Lau utilise a "network of queues" [AML13] in which queue behaviour and queue waiting time measurements are composed into a model, enabling the computation new estimates of queue waiting times for individual attractions within a theme park. Generally, mobility modelling is often used in urban environments to capture the flow of pedestrians or visitors and measure potential impact of changes in such environments. Similar to previous approaches, Pan et al. have used an agent-based human mobility model to understand the behaviour of crowds in evacuation scenarios [Pan+07]. Zhong et al. utilised a similar approach in the context of evacuation of train stations to understand the flow of people and evaluate whether travellers can be safely evacuated [Zho+08].

The use of simulation technology as a concept has previously also been explored in the context of digital signage deployments. Ostkamp and Kray have created visual simulations of displays situated in environments to explore potentiual locations for new display deployments and measure the perception of such deployments with "augmented panoramic imagery" [OK14].

3.4.3.4 Combination with Real-world Datasets

The core characteristic of our synthetic analytics system is the ability to bring together synthetic movement traces with real-world datasets. In our use case, the output of the mobility simulation consists of a list of display sightings per agent, enabling us to map a timestamp to each of the display sightings and combining these with Pheme's log of the content shown. More specifically, each display sighting consists of the following attributes: 1. the unique identifier of the agent, 2. display identifier and location, and 3. the timestamp of the display sighting. In order to retrieve the content shown at any particular time for any display on campus, we access Pheme via an API with the location identifier and timestamp and retrieve a set of metadata for the piece of content that was displayed at the given location and time. The metadata includes the name of the content, a unique content hash and, for Web-based content, the full URL. We expand the set of display sightings with the retrieved metadata from Pheme providing us with a foundation to compute new insights into the perception of content across the entire display network. In particular, we are able to consider and track individual synthetic viewers across multiple locations and displays without violating the viewers' privacy and enabling us to shift towards a *viewer-centric* analytics approach in the digital signage context.

3.4.3.5 Implementation

The synthetic analytics tool has been implemented in its entirety in Python (total $1,796$ lines of code). Whilst the spatial layout of the University campus has been colour-coded with standard image manipulation tools, the resulting encoded map was parsed using Pillow[5] and the extracted features were inserted into an undirected graph-based data structure provided by NetworkX. The encoded campus map (2119 by 5122 pixels) resulted in $11,225$ nodes ($6,827$ nodes representing buildings, $4,398$ nodes representing paths) and $29,539$ edges. Each node corresponds to a square on the map of 5 sqm. Models were implemented in the form of Python classes, whilst agents were modelled as instantiations and objects of these classes. Each mobility model class consists of a function that allows agents to return the movement decision in each iteration of the simulation. To retrieve display content logs from Pheme, we developed a dedicated export component that provides us with an API and the ability to query the content shown for a specific display location and timestamp. In order to improve the performance of the implementation by avoiding duplicate API calls, we implemented a simple caching mechanism in which we store returned values from Pheme regarding historical content shown in a local MySQL database—due to the nature of the data we do not request such information twice.

3.5 Summary

In this chapter, we described the importance of data collection for digital signage analytics. Concretely, we made the following four contributions.

1. We highlighted the potential for sharing analytics data sets and provided a framework for describing potential combinations of data.
2. We highlighted the importance of viewer-centric analytics in the context of open pervasive display networks.
3. We illustrated how digital signage analytics data can be collected and processed by techniques brought together from the Web analytics domain.
4. We designed three distinct approaches to collecting viewer mobility data: client-based tracking, infrastructure-based tracking and the use of synthetic mobility traces as a new approach to addressing privacy-related issues.

In the following chapter, we will provide examples of novel display analytics reports that utilise the datasets described as part of this chapter.

[5] Python Image Library

Chapter 4

Reporting

4.1 Overview

In this chapter, we explore novel analytics reports for the digital signage domain using the previously described opportunities for data collection as a foundation. Concretely, we follow the overall approach shown in Figure 4.1 and describe the opportunities particularly but not limited to viewer-centric analytics that can be produced by combining display-oriented data with viewer mobility patterns. We provide a set of example reports that have been created from the use of synthetic analytics (i.e. the use of synthetic mobility traces), and subsequently describe the extensions required to support analytics reports relating to display personalisation. We further describe the commonalities and applicability of digital signage analytics to Web analytics terminology, and demonstrate the potential of repurposing existing analytics services for other domains using Pheme.

We note that in the context of this chapter we focus purely on developing and describing novel analytics reports by applying existing information visualisation techniques, and describing the benefits and opportunities from different types of analytics reports for each of the stakeholder groups identified in the earlier chapters. We are not seeking to develop novel data visualisation techniques and algorithms for the extraction or detection of specific usage

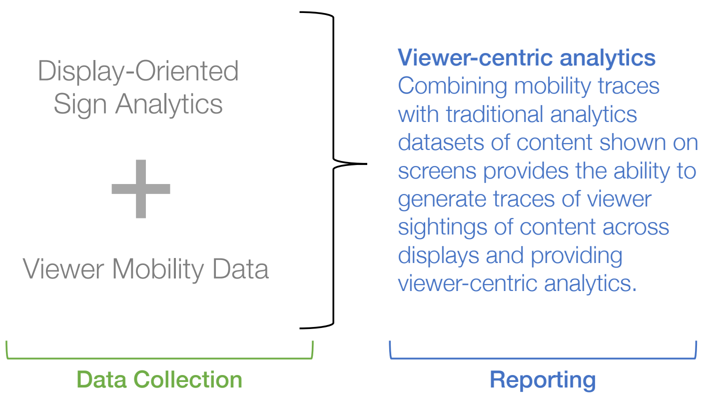

Figure 4.1: Using display-oriented analytics and viewer mobility data as a foundation, we now focus on novel viewer-centric analytics reports.

patterns. Additionally, we do not seek to conduct a comparison of different data collection techniques such as synthetic analytics and Tacita – instead, we purely focus on utilising data originating from the different systems to highlight the opportunities of novel analytics reports and insights that can be gained from such data.

Excerpts of this chapter have been published in the following peer-reviewed publications:

1. Mateusz Mikusz et al. "Repurposing Web Analytics to Support the IoT". in: *Computer* 48.9 (Sept. 2015), pp. 42–49. ISSN: 0018-9162. DOI: 10.1109/MC.2015.260. URL: http://doi.org/10.1109/MC.2015.260

2. Mateusz Mikusz et al. "Next Generation Physical Analytics for Digital Signage". In: *Proceedings of the 3rd International on Workshop on Physical Analytics*. WPA '16. Singapore, Singapore: ACM, 2016, pp. 19–24. ISBN: 978-1-4503-4328-2. DOI: 10.1145/2935651.2935658. URL: http://doi.acm.org/10.1145/2935651.2935658

3. Mateusz Mikusz, Sarah Clinch, and Nigel Davies. "Design Considerations for Multi-stakeholder Display Analytics". In: *Proceedings of the 6th ACM International Symposium on Pervasive Displays*. PerDis '17. Lugano, Switzerland: ACM, 2017, 18:1–18:10. ISBN: 978-1-4503-5045-7. DOI: 10.1145/3078810.3078830. URL: http://doi.acm.org/10.1145/3078810.3078830

4. Mateusz Mikusz et al. "Experiences of Mobile Personalisation of Pervasive Displays". In: *ACM Transactions on Computer-Human Interaction – TOCHI (in preparation)* (2018)

5. Mateusz Mikusz et al. "New Challenges in Saturated Displays Environments". In: *IEEE Pervasive Computing* (2018)

4.2 Analytics Based on Viewer Data

Traditional digital signage analytics are largely display and content focussed, providing only few insights into the experiences and interactions of passers-by with and across a display network. In this section, we explore the opportunities that emerge by capturing a comprehensive set of viewer mobility patterns and combining these with logs of content shown on displays for the creation of novel analytics reports. In addition we consider user interaction data captured through Tacita to create a novel set of retention analytics reports specific to support personalisable display networks. We note that the example reports and demonstrations of the usefulness of the insights are provided in the context of the Lancaster University deployment (initially described in Section 1.3, p. 6).

4.2.1 Datasets and Methodology

As previously described in Section 3.4 (Capturing Viewer Mobility Data, p. 65), a number of techniques can be used to capture viewer mobility data. In particular, we consider datasets that are captured through mobility models (synthetic mobility traces), and viewer-based tracking (originating from Tacita). We note that the use of infrastructure-based data capture techniques was not possible due to the lack of an appropriate Wi-Fi location tracking system at Lancaster.

4.2.1.1 Synthetic Analytics

We consider the use of synthetic movement traces created using the synthetic analytics approach first introduced in Section 3.4.3 (Synthetic Analytics, p. 76) utilising the following mobility models: *On-Campus Student*, *Off-Campus Student* and *Random Building Navigator* (see Section 3.4.3.2, p. 78 for more details).

Using these models, we executed the simulation with $2,000$ agent instances for each mobility model simultaneously resulting in a total of $6,000$ agents constantly moving across the spatial model of the University campus. We simulated a specific time period of 62 days (1 October 2015 until 1 December 2015) and combined the computed mobility traces and display sightings resulting from the simulation with the logs of content played from the identical time period that were captured using Pheme (the dataset captured through Pheme was initially described in Section 3.3, p. 59). The resulting dataset therefore consists of a combination of both real and synthetic analytics data and provides a set of display sightings and a log of content seen for each instance of an agent. This dataset enables us to create analytical insights for individual content items across the signage network and the experiences of viewers and passers-by – without the limitation to single displays or isolated spatial areas of the deployment. These complex signage analytics insights would have otherwise required the deployment of comprehensive tracking technology of individuals – going beyond what current state-of-the-art analytics tools (e.g. face recognition software) are able to offer in the digital signage domain.

We note that accuracy of reports generated using synthetic analytics are highly dependent on the quality of the underlying mobility models. In future, movement and mobility traces collected through various tracking technology could be used to inform the design of corresponding models and ultimately improve the quality of the simulation. We recognise that our approach is limited due to the use of simplistic mobility models and the lack of the consideration of contextual events. For example, lecture timetables, bus schedules and automatically collected room occupancy metrics (e.g. through the use of attendance monitoring software) across the University campus could be used to inform the design of mobility models.

A complete description of the dataset captured through Pheme is provided in Section 6.2 (Pheme: Display-oriented Data Collection, p. 125) as part of the evaluation of the system.

4.2.1.2 Tacita

In addition to synthetic synthetic analytics as a source for capturing and generating mobility traces, we also consider data captured from Tacita. The dataset resulting from Tacita has been initially described in Section 3.4.1.4 (Opportunities for Data Collection, p. 71). In particular, we are able to capture the following insights: configuration parameters of Trusted Content

Providers (i.e. the viewers' personalisation preferences), display sightings of viewers (i.e. viewers are detected in proximity of displays), and the Trusted Content Provider that was shown on the display when the viewers were detected in proximity of the display. We note that we have conducted a long-term, large-scale trial that will be described in more detail in Section 6.2 (Pheme: Display-oriented Data Collection, p. 125). The dataset emerging from Tacita is used to create similar types of analytics reports and, in addition, a set of reports that are specific to display networks that support the delivery of personalised content.

4.2.1.3 Comparing the Applicability of Synthetic Analytics and Tacita

Both synthetic analytics and Tacita have advantages and disadvantages as potential sources for viewer mobility data. Whilst both approaches can be used to underpin the same categories of analytics reports (as we illustrate in the subsequent sections), the data capture process differs significantly and each of the approaches are applicable in certain contexts.

The synthetic analytics approach is particularly practical if certain technical, legal or ethical constraints prevent space owners from applying viewer- or infrastructure-based tracking mechanisms. Synthetic analytics relies purely on appropriate viewer mobility models and does not require additional hardware to be deployed. If, for example, the deployment of Bluetooth Low Energy beacons is impractical, synthetic analytics can be used as a solution in order to gain insights on interactions and movement patterns. Whilst we applied synthetic analytics to capture display sightings of viewers, additional information can be encoded in the model to capture a broader set of insights – such as more broad behaviour and navigation patterns that go beyond simple display sightings and provide further insights on peoples' interactions across a space.

In contrast, Tacita provides insights that are potentially of higher accuracy as display sightings are based on Bluetooth Low Energy sightings instead of a synthetic model of user mobility. Tacita also presumes that viewers (or users of the display personalisation system) have explicitly opted in to the required display proximity tracking – leaving it up to the user to decide whether to contribute data to the system (an opt-in and opt-out in the synthetic analytics approach does not make sense). The use of Tacita, or more generally viewer-based analytics tracking approaches, allows for additional features such as personalisation as described in previous chapter to provide a visible benefit to the user in exchange for location tracking. We note, however, that the deployment of a system such as Tacita requires substantial investment from display and space owners: displays need to be equipped with appropriate hardware infrastructure (e.g. Bluetooth Low Energy beacons), and a corresponding mobile phone application and backend components are required to allow for capture and reporting of user locations in the space.

4.2.2 Effectiveness of Displays

The access to mobility traces of individual people (in the case of synthetic analytics, individual agents; and for Tacita, display sightings of users) enables us to understand the effectiveness

Table 4.1: Ranking of public displays at Lancaster University based on synthetic analytics.

Ranking	Display Location	Total Views
1	Faraday Centre	270,625
2	Faraday Left	259,685
3	Faraday Right	250,820
4	Alex Square West	201,744
5	Nuffield External	152,750
...	...	...
14	ISS	11,521
15	Student Services	7,528
16	Human Resources	1,763
17	Cartmel College (2)	568
18	Cartmel College (1)	463

and visibility of displays deployed within a space based on their geographic location and numbers of passers-by. Knowing the identities of each individual agent and their associated display sightings enables us to calculate both the total number of viewers and the proportion of unique viewers that are passing by displays across the space throughout the total length of the simulated time period.

4.2.2.1 Reports Based on Synthetic Analytics

Firstly, we utilise the *Random Building Navigator* as a baseline to quantify the visibility of displays. The mean number of display views across the entire display network is 73,299 (SD: 81,067). The numbers reported below are based on the combination of all three mobility models (*On-Campus Student*, *Off-Campus Student* and *Random Building Navigator*).

Based on the number of total views for individual displays, we are able to calculate the highest and lowest ranked displays, i.e. displays with the highest and lowest number of views (in this case, we count repeating views of the same agent). Table 4.1 provides an overview of the ranking of displays. Among the highest ranked displays, we find displays located outside the Faraday lecture theatres, a building with three lecture theatres and one of the highest used at the university explaining the high volume of agents passing by causing the highest view counts. Additionally, entrances and exits of this building are often used to pass through to adjacent buildings causing a high transit traffic. The fourth and fifth highest ranked displays are Alexandra Square West and Nuffield, both located outdoors along main pathways on campus often used by students to pass between colleges and lectures. The poorest ranked displays include displays located in the outskirts of the campus in student accommodation (Cartmel College displays), and displays located in buildings mainly used by members of staff (Human Resources and ISS) which typically have a very little number of people passing by, leading to very low view counts for displays located in these buildings.

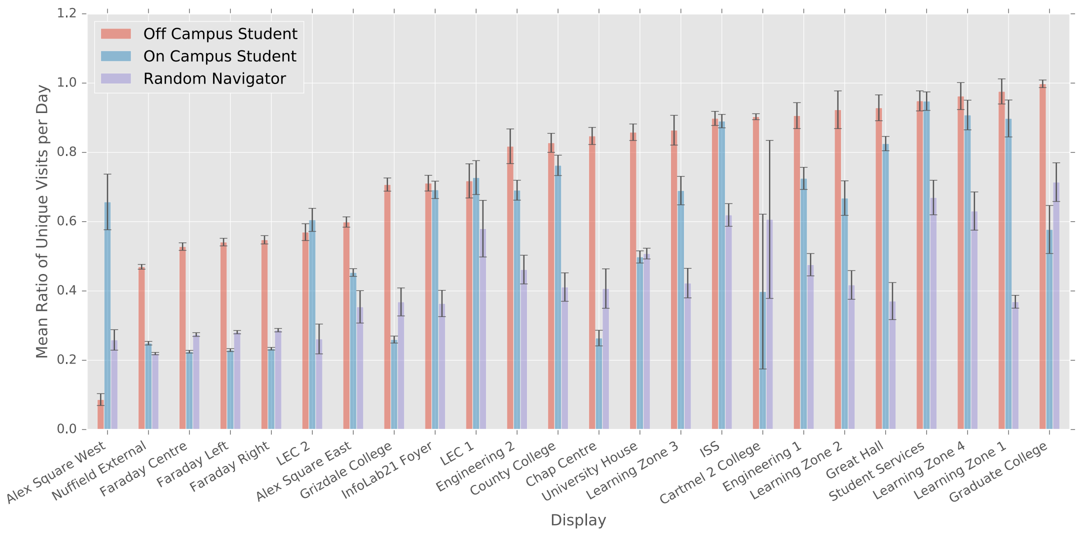

Figure 4.2: Mean ratio of unique visitors to total number of visits for each class of mobility model (ordered by the ratio for the *Off-Campus Student*) for each display deployment. Low ratios indicate displays that have a greater set of reoccurring visitors whilst high values represent displays with a greater set of unique visitors (initially published in [Mik+16]).

While reports on viewer sightings could be achieved using conventional video analytics (alongside with the cost of equipping displays with cameras), due to the ability to distinguish between new and returning viewers, we were able to provide new insights including the mean ratio of unique visitors to the total number of views for each display as shown in Figure 4.2. Based on the analysis, we find displays with a lower ratio of unique viewers (i.e. displays with a high proportion of returning viewers) are dominated by displays located outdoors along the main pathway on campus (e.g. LEC and Nuffield displays), and near the main bus station (Alexandra Square West). Similar to the previous findings based on the number of total views, displays with a high ratio of unique viewers are located on the outskirts of the campus and in student accommodations (e.g. Graduate College), and in central learning spaces (e.g. Learning Zone).

Such insights into the ratio of unique viewers can be used for optimising the scheduling of content per display, e.g. for displays with a higher ratio of unique viewers it might be beneficial to ensure that only a small set of content items is scheduled to show to ensure that passers-by have a chance to notice and see important pieces of content whilst on displays with a high ratio of returning visitors, passers-by will have more opportunities to see a particular piece of content.

4.2.2.2 Reports Based on Tacita

The data collected as part of Tacita enables us to create a comparable set of analytics reports to those created using synthetic analytics. Whilst we are unable to capture actual display sightings, we were able to capture Tacita users entering the proximity of a display (resulting in a set of personalisation requests). Based on this dataset, we created a ranking of pervasive displays based on the total number of personalisation requests received throughout the entire

Table 4.2: Ranking of public displays at Lancaster University based on display personalisation requests from Tacita throughout the lifetime of the service (initially published in [Mik+18d]).

Ranking	Display Location	Total Requests
1	Library (C-Floor)	968
2	Library (B-Floor)	950
3	InfoLab Foyer	844
4	Library (A-Floor)	786
5	Learning Zone 2	622
...	...	... ...
36	Bowland College	10
37	Bowland Main (B-Floor)	10
38	Pendle College	10
39	Welcome Centre (2)	10
40	Welcome Centre (1)	6

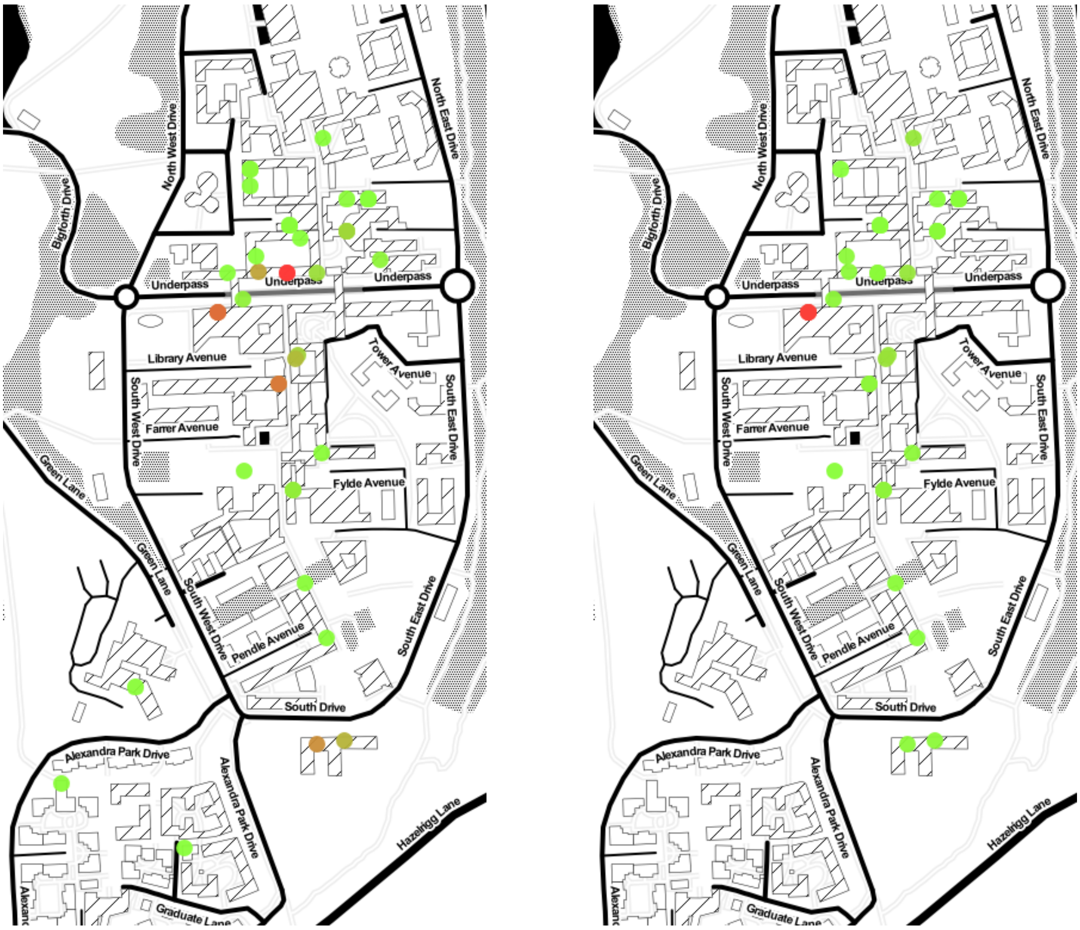

(a) During lecture days and times. **(b)** During non lecture days and times.

Figure 4.3: Heatmap of proportionally popular displays on which personalised content was requested (initially published in [Mik+18d]).

deployment. Table 4.2 shows that the list of highest and least ranked displays is somewhat comparable to the display rankings created in the synthetic analytics approach. Among the highest ranked displays, we see displays located across various floors in the University library (Library A-Floor, Library B-Floor and Library C-Floor with 968, 950 and 786 personalisation requests respectively) and the Learning Zone (622). These displays are located in areas highly frequented by students especially during typical exam and study times with high numbers of personalisation requests to be expected. The third-highest ranked display is InfoLab Foyer (844 personalisation requests). This display is located in the foyer of a departmental building with a high frequency of visitors and staff passing by and is part of the original set of personalisable displays. Among the least ranked displays, we find Bowland College, Bowland Main and Pendle College (each 10 views respectively). These reports confirm previous findings using the synthetic analytics approach where displays located in University colleges retrieved only limited exposure levels. The lowest ranked displays are Welcome Centre 1 and 2 (10 and 6 requests respectively) located in a remote part of the University campus typically only used by campus visitors who have not been introduced to Tacita.

Using the same dataset as a basis, we created a set of heatmaps to further visualise common areas on campus with a high and low numbers of display personalisation requests during specific times. Figure 4.3a shows a heatmap of personalisation requests during typical lecture times only (weekdays between 8am and 6pm excluding lunch breaks between 12pm and 1pm), whilst Figure 4.3b shows personalisation requests during lecture-free times (weekends, and weekdays before 8am, during core lunch times, and after 6pm). Comparing both heatmaps we note that the variation is only marginal. During typical lecture times, we see an increased number of personalisation requests in locations near lecture theatres, whilst in lecture-free times the number of requests in typical study zones (e.g. the Learning Zone and library) increases proportionally to the total number of requests captured in these time frames.

4.2.2.3 Stakeholder Analysis

The reports outlined above provide the following set of benefits to individual stakeholder groups.

Display Owner Display owners can utilise the reports described above to understand the level of visibility (measured in number of total views and unique viewers) of their displays. This information can help stakeholders evaluate existing display locations, and select appropriate types of content for displays. For example, for displays with a high proportion of recurring viewers it may be more appropriate to schedule a broader set of content items to offer a higher variety to a smaller viewer group. However, displays visited by a large number of different viewers (and having a high volume of views) may only need a smaller set of content items to maximise the visibility of individual content. In addition, reports such as those described above can be used to measure the impact of content scheduling changes applied by the display owner in order to understand the potential influence of content to the behaviour of viewers.

Space Owner The insights described above are useful for space owners to evaluate existing display location placements. The number of total views and unique viewers can be used to determine displays with large or small viewer sets – e.g. to remove displays with only a very low visibility or prioritise maintenance work based on the number of views. Additionally, viewer numbers of displays can also be used to understand the general flow of individuals and crowds throughout their space in order to, for example, optimise the flow in their space. Using synthetic analytics, future display locations can be simulated and evaluated (e.g. by measuring the impact on the flow of individuals) in order to determine whether a location offers benefits to display owners.

Content Provider The effectiveness of displays measured in views and unique viewers provides detailed insights on the potential reach of content. Content providers can decide based on the number of views which displays may be more or less appropriate for their content. In addition, similarly to display owners who can experiment with different types of content, content providers can measure the impact of different ways in which the same content is delivered (e.g. different formatting) in numbers of views.

Viewer Whilst reports regarding the effectiveness are primarily targeted at administrators and display providers, viewers are able to indirectly benefit from such reports by an improved quality of the display deployment (e.g. better display placements) and more appropriate content scheduling.

4.2.3 Network Visibility of Content

The knowledge of display sightings (both based on synthetic analytics and Tacita) enables us to create reports regarding the content visibility in addition to the 'content impressions' metric commonly used in signage and Web analytics. We note that in contrast to the Web domain, the fact that a content item was shown on a display does not mean that that the content was viewed, e.g. due to the lack of audience or viewers not paying attention to the display. The knowledge of display sightings, however, allows us to create more accurate estimations of content views. Additionally, we can disconnect content from displays and consider the visibility of content across the entire display network instead of a 'per display' basis.

4.2.3.1 Reports Based on Synthetic Analytics

Drawing on the synthetic analytics approach and using the *Random Building Navigator* as a baseline, we computed a mean number per content item per day of 217 total views (SD=507) and 116 unique views (SD=190). The most viewed content items include content that has been scheduled onto one of the most viewed displays on campus exclusively (Nuffield External) and content that is scheduled across the university campus. Content with the least views include, not unexpectedly, content that has been scheduled on single displays only (e.g. ISS).

Understanding in more detail how viewers experience content and view the same piece of content multiple times is particularly interesting for content that has been scheduled to show

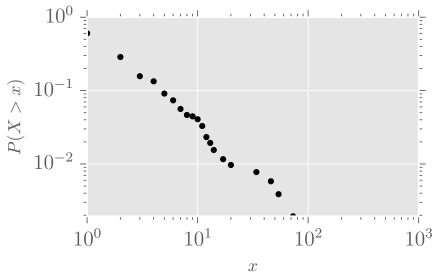

(a) Random Building Navigator.

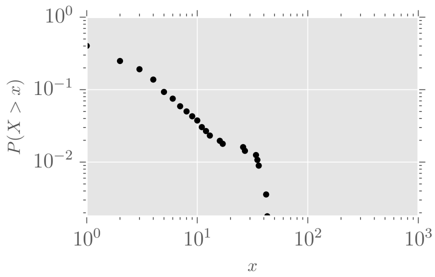

(b) On-Campus Student.

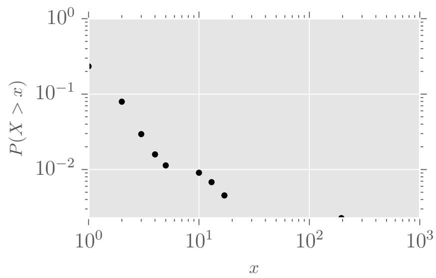

(c) Off-Campus Student.

Figure 4.4: Cumulative distribution function of frequencies of content views for the three synthetic analytics mobility models (initially published in [Mik+16]).

on multiple displays of the University campus. We can create reports that give us insights into the frequency of views of content items seen repeatedly by the same user. Content providers can benefit from such insights to optimise the design and distribution of content. For example, content providers might wish to design their content appropriately if it is viewed repeatedly, whilst other content providers might prefer to reduce the visibility of their content items across a signage network. To retrieve insights into repetitive views for the same content items, we have created a cumulative distribution function of the frequencies of content views (Figure 4.4). Whilst the *Random Building Navigator* (Figure 4.4a) and *On-Campus Student* (Figure 4.4b) view the same pieces of content with similar frequencies, the *Off-Campus Student* (Figure 4.4c) views items with a lower frequency with the exception of a single piece of content viewed with a noticeably higher frequency. This is likely due to the mobility patterns of the model: students arrive at the central bus stop on campus which is located near the Alexandra Square West display, typically only showing real-time bus timetables as the single piece of content, and is therefore seen by most students arriving and leaving from this bus stop.

Using our simulated location traces in combination with content views, we can additionally create reports to estimate the duration viewers spend looking at displays. We note, however, that this statistic is highly approximate due to a number of factors. Firstly, the views are an approximation themselves and based on probability and mobility models. Secondly, to compute the mean durations viewers spend glancing at displays, we make use of previous work focussing on measuring the mean time passers-by spend glancing at displays. Dalton, Collins, and Marshall [DCM15], for example, suggest based on an eye-tracking study that passers-by glance at displays for a mean duration of 0.318 seconds (SD: 0.261) whilst work

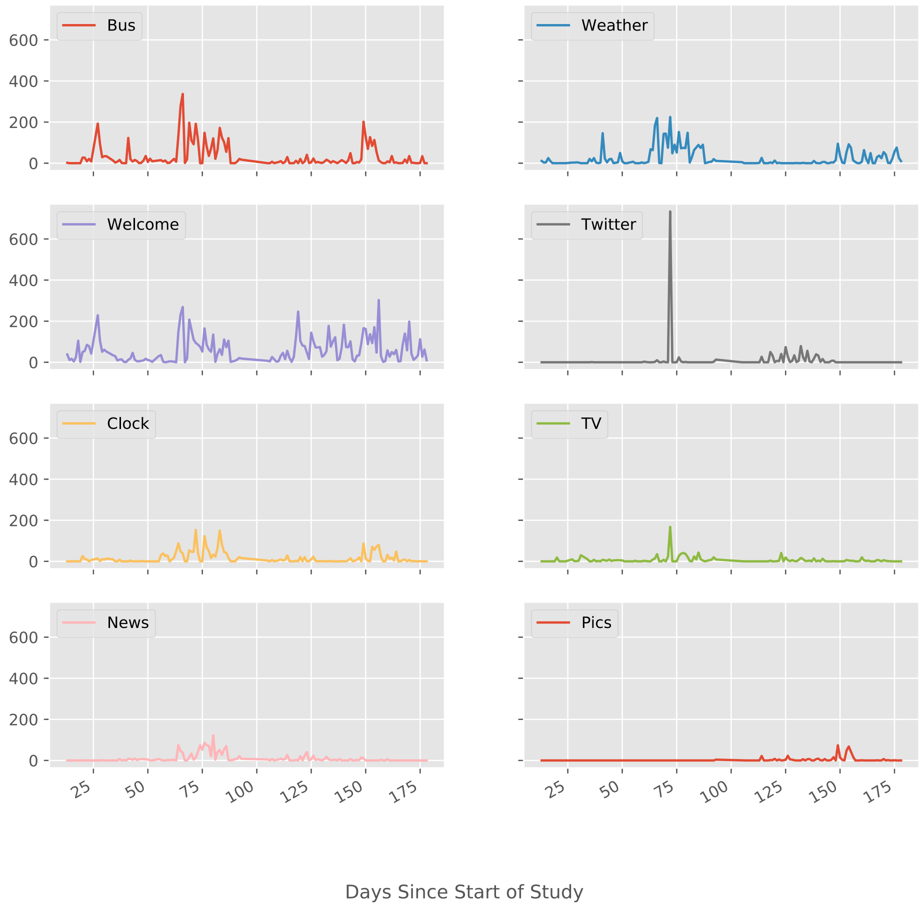

Figure 4.5: Number of total requests per Tacita application per day (initially published in [Mik+18d]).

by Huang, Koster, and Borchers [HKB08] indicates based on observational studies a glance period of 1-2 seconds. Using both figures as the lower and upper bounds, we are able to compute mean view times between 1.15 and 5.43 minutes per content item per day aggregated across all agents. The mean total view duration per content items equates to 20.33-95.88 minutes.

4.2.3.2 Reports Based on Tacita

Tacita data can also be used to generate the same types of analytics as described above (Section 4.2.3.1). However, in Tacita users request personalisable content (Trusted Content Providers) automatically when their presence is detected in proximity to a display. Using these display sightings, we can produce new analytics that consider the visibility of such personalisable content requests and therefore use it to provide insights about the usage of Tacita applications. For example, we can consider the total number of requests of personalisable applications issued by the users' mobile devices when a display was detected in proximity to understand better which applications are more popular. Figure 4.5 shows the total number of requests per application per day across the University campus. Such a report can be used by both display owners and content providers to better understand the popularity of individual

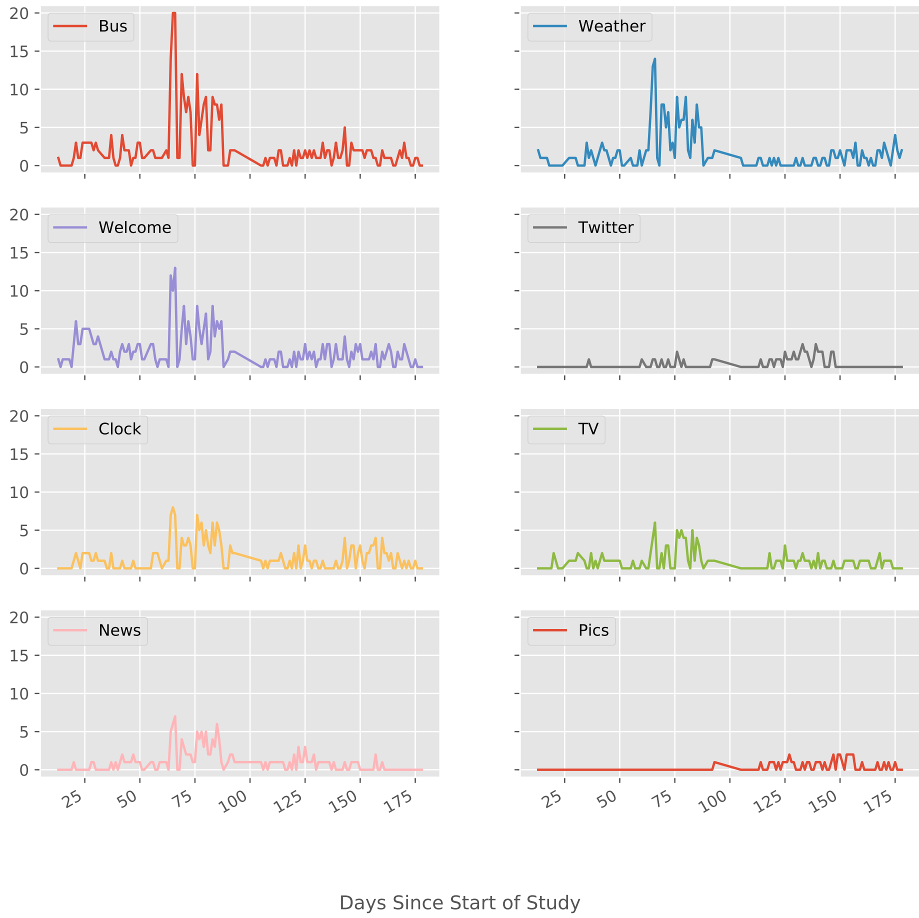

Figure 4.6: Number of unique users per Tacita application per day (initially published in [Mik+18d]).

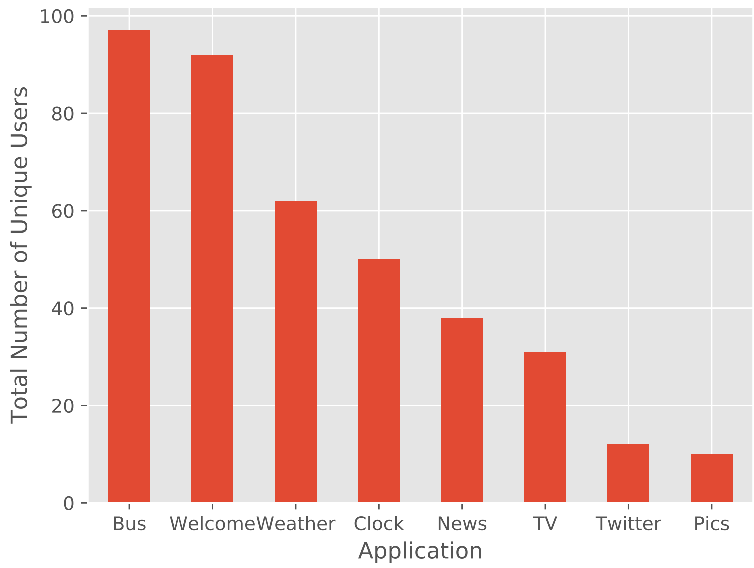

Figure 4.7: Total number of unique users per Tacita application (initially published in [Mik+18d]).

content items and applications shown on the display network without the focus on individual displays or users. Using content requests as a metric, the most popular applications include Bus Timetables (personalised travel), Weather and Clock, whilst we find Twitter (personalised social media) among the least requested displays. The spread and popularity of display applications may reflect the property of *public* displays: users prefer not to retrieve highly personalised content such as a newsfeed from their social media.

Whilst Tacita only captures anonymised user interaction and content requests, we can use the anonymised unique user identifiers to further narrow down the reports and report *unique user* metrics similar to what can be found in typical Web analytics. For example, Figure 4.6 shows the number of unique users per day for each of the supported application. In contrast to reporting the total number of requests, the reports using unique users enables us to better understand the popularity of individual applications disconnected from the frequency or mobility patterns of users that may ultimately lead to high or low requests.

Moving away from reports regarding individual days, we are able to confirm these findings by creating reports considering the *total number of unique users*, i.e. users who have requested an application at least once in the lifetime of the deployment. Figure 4.7 provides a bar chart ordered in descending order by the application with the highest number of unique users. The most popular application includes Bus Timetables, followed by the Welcome application. This result indicates that whilst a large number of users configures and activates other application, they choose not to deactivate the preconfigured Welcome application before the first appearance at a display.

4.2.3.3 Stakeholder Analysis

The reports outlined above provide the following set of benefits to individual stakeholder groups.

Display Owner Display owners can gain a detailed understanding in the interaction patterns of users with their displays and the types of content viewers have been interested in across the entirety of the display network. In the context of Tacita, display owners get insights into the types of personalisable applications that have received the most interest among viewers.

Space Owner Space owners typically have only limited control over the displays deployed in their space. Providing insights into the visibility of content across the entire network (or across spaces owned by individual stakeholders) can provide crucial insights on the types of content that is shown in the space, and the popularity of different categories of content.

Content Provider The knowledge on the duration individuals viewed content items, and how often content was viewed across the entire display network can be valuable to content providers. For example, durations of content seen can be used to optimise the amount of information presented on a screen (e.g. include less content on the display if viewers see content less often or for shorter periods of time overall). Furthermore, content providers can use these analytics reports as confirmation that content has been scheduled and seen across a display network. In the context of personalisable applications, analytics reports described above can be used for monitoring purposes in order to understand whether certain applications receive viewer location information and are successfully scheduled onto displays.

Viewer Viewers benefit from content-related reports indirectly by improved content scheduling cycles and, more generally, an improved quality of the content presented. Individual content can be better tailored towards viewers if, for example, content that has been seen less frequently in relation to other content is prioritised across the display network to improve its visibility.

4.2.4 Viewer-centric Analytics

We explored the creation of reports that capture the viewer experience and visibility of content across the entire signage network. Being able to track individual display views and associate these with individuals enables us to approximate the number of unique content items individuals see on average. In the context of Tacita, we are able to create detailed analytics regarding the kinds of personalisable applications users experience throughout a day on average.

Table 4.3: Aggregated count of unique and total content views per day per viewer for each mobility model (initially published in [Mik+16]).

	unique views			total views		
Mobility Model	mean	median	SD	mean	median	SD
Off-Campus Student	4.49	4.00	2.70	13.76	14.00	8.16
On-Campus Student	8.77	7.00	5.91	16.01	10.00	20.06
Random Navigator	8.14	7.00	5.36	16.02	11.00	17.75

4.2.4.1 Reports Based on Synthetic Analytics

Table 4.3 provides an overview of the median and mean number of total and unique content items each of the agents in each of the mobility models view throughout the simulation within a day. For example, the *Off-Campus Student* sees a mean of 4 different pieces of content in a day, whilst the *On-Campus Student* and *Random Building Navigator* yield similar results and see approximately 8-9 unique pieces of content throughout a day. However, both the *On-Campus Student* and *Random Building Navigator* have higher content views (mean: 16) within a day whilst the *Off-Campus Student* is generally exposed to less content (mean: 13.76). Similar to the previous reports, the number of content items experienced again reflects the specific constraints of the individual mobility models. For example, the agents of the *Off-Campus Student* model arrive and depart from a central part of campus. In addition, the high standard deviations for mean and median content views for all mobility models are likely a result of the differences in content schedules across displays on the University campus. Some displays, for example, consist of a high mix of different content items reducing the probability to see the same content twice whilst other displays show a smaller mix of content that is also distributed to other displays on campus.

Using similar datasets, we are able to report *content repetition* metrics, i.e. describing the number of times a single viewer sees the same content multiple times. For the *On-Campus Student*, *Off-Campus Student* and *Random Building Navigator* models, we are able to report that viewers see the same piece of content on average 86.77, 54.63 and 94.86 times respectively (SD: 22.82, 12.47, 26.02) throughout the simulated time period. These numbers highly depend on the content schedule in pace during the simulated time period, in this case indicating infrequent changes in the set of scheduled content items.

Similar to reporting the durations of looking at individual content items, we can use the collected content sightings to create reports about the (mean) durations viewers spend glancing at content items throughout the entire deployment and simulation. Using similar estimates of average durations viewers spend looking at displays, the *On-Campus Student*, *Off-Campus Student* and *Random Building Navigator* spend a total of 4.52-21.33, 5.26-24.82 and 5.26-24.83 minutes respectively looking at content in the simulated time period. During a day, viewers spend only 5-25 seconds viewing content. Divided by the number of viewed content items, this equates to less than 0.714 seconds for the *On-Campus Student* and *Random Building Navigator* models and 1.25 seconds for the *Off-Campus Student* model.

Table 4.4: User-centric statistics for the use of Tacita (initially published in [Mik+18d]); total number of users: 147.

Metric	mean	median	SD
Number of apps per user	2.68	2.0	1.52
Number of apps per user (excl. Welcome)	2.2	2.0	1.33
Number of total requests per user per day	134.8	55.0	208.08
Number of unique users issuing requests per day	4.59	3.0	4.51

4.2.4.2 Reports Based on Tacita

Once again, Tacita data can be used as an alternative to synthetic mobility traces for the creation of reports. In addition, we can provide reports regarding the usage patterns of individuals. Table 4.4 provides an overview of key statistics and metrics around the Tacita service from a total of 147 users throughout the lifespan of the deployment. Tacita users typically activated and requested a mean of 2.68 applications (Mdn: 2.0, SD: 1.52) out of the total of eight available applications. If we exclude the pre-activated Welcome application, Tacita users requested a mean of 2.2 applications (Mdn: 2.0, SD: 1.33) indicating that Tacita users typically configure and activate applications in addition to the default settings.

To consider the volume of personalisation requests issued by users per day, we created reports consisting of the *number of personalised content requests users typically issue within a day*, and the *number of unique users issuing at least one request per day*. The mean number of total requests issued on a daily basis is 134.8 (Mdn: 55.0, SD: 208.08). The high standard deviation is likely due to the fluctuation of viewers on campus, e.g. weekends and holidays in which the number of requests and users drops significantly. Further, we believe a small set of users request high amounts of data whilst a separate set of users only issues a low number of personalisation requests – either due to their navigation patterns only rarely leading by displays or by their preferred usage of the service. For example, some users might choose to disable the background location tracking of Tacita but prefer a more explicit usage model in which personalisable applications are requested by explicitly opening the mobile phone application. Over the whole deployment, we computed a mean of 4.59 unique users requesting personalised content at least once each day (Mdn: 3.0, SD: 4.51). Similarly to the previous report on total numbers of requests, the relatively high standard deviation reflects the nature of the long-running deployment and the naturally low numbers of requests outside of term times and during holidays.

Overall these reports enable us to consider that most users are interested in a very low number of personalisable applications (Mdn: 2.0, SD: 1.52). Such low numbers of applications indicate that personalisable display deployments will have little issues with competing applications for single users. However, the competition among different application providers is likely to be high with users preferring to choose a low number of applications only.

4.2.4.3 Stakeholder Analysis

The reports outlined above provide the following set of benefits to individual stakeholder groups.

Display Owner Display owners can receive detailed insights into the viewer experience throughout the entire display network measured by the number of (unique) displays and content viewers typically consume throughout a day. Such numbers can be particularly useful if display owners wish to explore alternative content scheduling algorithms to understand the impact on the viewer experience. For example, if the total number of content views is significantly higher than the number of unique content items seen, display owners may adjust the content scheduling in order to give individuals the opportunity to see varying content instead of the same content multiple times.

Space Owner Space owners can use the reports to understand how viewer experience is influenced by displays deployed in their space and the exposure of viewers to content. Such insights can be used to inform the decision on potentially allocating more or less space for display deployments if the content exposure is too low or too high respectively.

Content Provider Understanding the frequency of display and content views can provide valuable insights for content providers regarding the number of opportunities individuals typically have to consume the content provided. Such understanding can then inform the content delivery decisions (e.g. in the context of personalisable applications) or help shape the design of content items and the level of information present. For example, if viewers typically only see a small number of displays and content items throughout a display network, content providers can design the content accordingly by including all relevant information in single content items.

Viewer Viewers provided with access to the analytics reports described above can better understand their own exposure to public displays and the content displayed across a day and display network. In addition, viewers may benefit from an improved quality of content scheduling algorithms and better designed content.

4.2.5 Display Personalisation Retention Analytics

The dataset captured through Tacita can be used to create an additional set of analytics reports specific to describing viewer engagement and interactions with Tacita. In contrast to analytics reports captured for purely commercial- and advertisement-based public displays and billboards, the requirements for analytics reports in the context of Tacita are around revealing and visualising insights into the viewer behaviour and usage of personalisable display applications such as the ways in which users configure content and the spatial distributions of requests for personalised content.

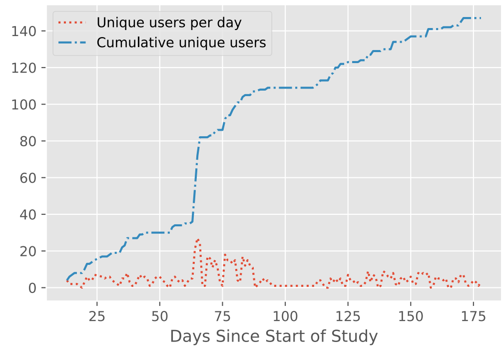

Figure 4.8: Growth of Tacita users throughout the deployment (blue) and unique number of users per day (red) (initially published in [Mik+18d]).

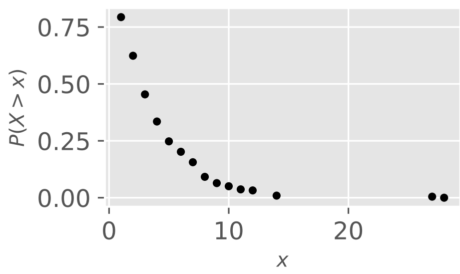

Figure 4.9: Cumulative distribution function of frequency of revisiting the configuration pages of individual Trusted Content Providers (initially published in [Mik+18d]).

4.2.5.1 Usage and Interactions

The assignment of a globally unique identifier to requests originating from Tacita users enabled us to recognise recurring users and created viewer-centric usage reports throughout the lifetime of the deployment. This included insights into the usage and interactions with Tacita including the reporting of *unique users* (Figure 4.8).

Tacita is characterised by a set of unique interaction patterns that differ to these found across mobile phone applications and traditional, interactive public displays. For example, unlike interactions with many mobile phone applications, Tacita users are only required to activate and configure the system once on their mobile devices. While passing by displays that support display personalisation, users implicitly interact with displays and applications simply by their phones detecting proximate displays and requesting applications that users have previously activated – users are not required to actively launch Tacita on their mobile device for further use.

Considering these unique interaction characteristics, we have created a set of reports that allow display owners and content providers to better understand the usage patterns and explicit interactions with configuration pages. Besides standard Web analytics that can be captured on the Web-based configuration pages, we considered each time users accessed and changed their preferences and calculated a cumulative distribution function (Figure 4.9). The figure visualises the proportion of users who revisit the configuration pages a number of times – expressing, in this case, that the majority of Tacita users configure their applications once and only a very small proportion of users revisit the page again to adjust configuration parameters. Such insights into interaction patterns can be crucial for the design of personalisable applications. For example, low frequencies of revisiting configuration pages may emphasise the importance of creating a good user experience as users are likely not to revisit configuration pages. We note that due to the nature of the configuration pages, user interactions within these pages, e.g. button clicks, scrolling behaviour and configuration parameters, can be captured and reported using standard Web analytics techniques.

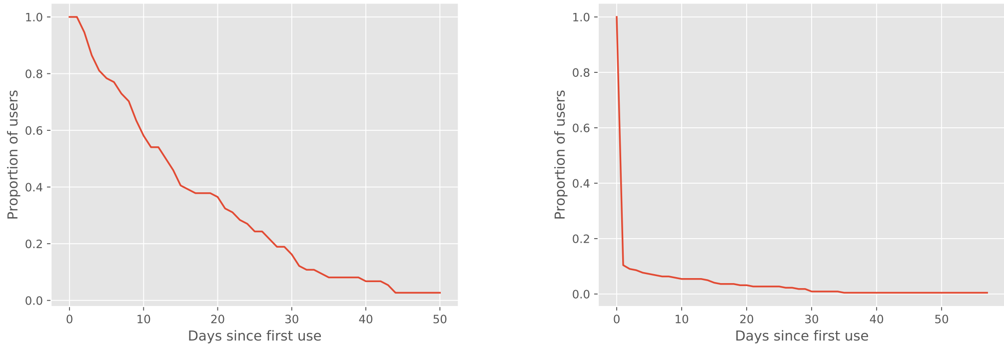

(a) User retention report based on display personalisation requests .

(b) User retention report based on configuration changes of personalisable applications.

Figure 4.10: Tacita user retention reports with a *per-day* granularity (initially published in [Mik+18d]).

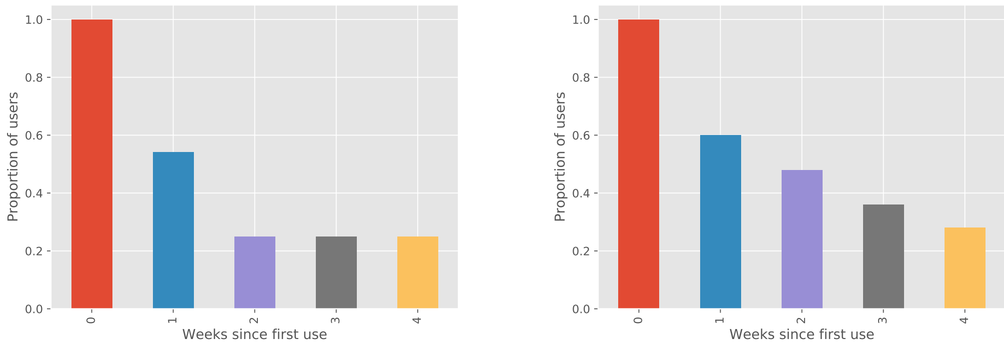

(a) User retention during winter term.

(b) User retention during spring term.

Figure 4.11: Tacita user retention reports with a *per-week* granularity (initially published in [Mik+18d]).

4.2.5.2 Retention Rates

A common way to report on the loyalty of customers in retail or users in a mobile phone setting are *retention rates* [RZ93]. Such retention rates can be used to describe the success of mobile phone applications: if users are still accessing a mobile phone application after multiple days, weeks or months it is considered as an indication for a successful application [Pel+18]. In the context of personalisable public displays, similar reports can be created to describe the lifespan of users requesting personalisable content on the display network. Whilst in mobile phone usage statistics retention rates are computed by considering the time span between the first and last time a user has opened an application, the metric can be adapted to our use case by considering the first and last time viewers requested personalised content on displays or explicitly visited the configuration page of a personalisable application.

To demonstrate the potential of the insights that can be gained from such retention reports, we explored two different types of reports with different granularities. Firstly, we consider retention reports on a per-day granularity showing the proportions of the activity lifespans based on the days counted between the first and last captured requests of viewers and between the first and last interaction with application configuration pages (Figure 4.10). Secondly, due to the nature of display personalisation systems (i.e. viewers are required to pass by supported displays in order to be considered in the reports), we additionally created reports with a lower granularity on a per-week basis accounting for the fact that some viewers could be still using the system but happen to not pass by a supported display (Figure 4.11). Figure 4.10a indicates a relatively high retention of continuous personalisation requests: after ten days, over 50% of users continued issuing personalisation requests implicitly by walking by displays indicating that Tacita has not been removed or disabled. In line with the cumulative distribution function of frequencies of revisiting configuration pages shown in Figure 4.9, the retention rate for configuration changes shown in Figure 4.10b shows a low retention rate after the first day indicating that users very rarely revisit configuration pages to change their preferences. Considering retention figures on a finer time granularity (Figure 4.11), we observe that over 50% of viewers are still successfully requesting personalised content one week after beginning to use the system – a metric that can be observed both during winter and spring terms (Figures 4.11a and 4.11b respectively). We further observe that whilst the majority of viewers cannot be observed after two weeks of usage, the proportion of 'long-term' users remains stable across weeks 2, 3 and 4.

4.2.5.3 Stakeholder Analysis

The reports outlined above provide the following set of benefits to individual stakeholder groups.

Display Owner Display owners can utilise retention rates to determine the perceived utility of the content (or personalisable applications) offered to the viewers. Very low retention rates of certain categories of applications may suggest that the overall offering of personalisable applications (and content) can be improved.

Space Owner Space owners gain insights into the popularity and uptake of novel technology deployed in their space. In the context of Tacita, for example, space owners would likely have been involved in the installation, deployment and advertisement of the application to visitors of the space. The use of retention rates can help space owners understand their 'return on investment' and potentially have an influence on the future deployment of technology for improving the experience of visitors or customers of the space.

Content Provider Retention rates in the context of digital signs can be used to inform the design of personalisable display applications and the content offered. For example, we noted very low retention rates of users accessing configuration pages of personalisable applications suggesting that users of such a system only configure the application once or very few times. Such usage patterns may introduce a set of system design requirements that may impact the design and development of applications.

Viewer Viewers can utilise retention reports to learn more about their usage patterns of personalisable applications and better understand why certain content was visible on individual displays when passing by.

4.2.5.4 Limitations

We note that computing and reporting retention rates in this context is subject to a set of limitations. In particular, whilst for mobile phone applications users actively interact with an application (i.e. at a minimum open the application), such user behaviour cannot be directly transferred into our use case: once Tacita applications have been configured and activated viewers are not required to actively engage with the mobile phone application or the displays. Instead, simply walking by supported displays is sufficient to automatically be detected in the proximity of the display and request the activated set of applications. Equally, if viewers do not pass by a display content requests are not issued but viewers could still be active users of the system. Consequently, assuming that viewers stopped using the service after the last content request does not take into account other potential reasons for viewers not requesting personalised content such as viewers choosing different walking paths not leading past displays or simply deactivating Bluetooth or other required functions on their mobile phones while still considering themselves as a user of the display personalisation system. These limitations are particularly important to consider when interpreting trends in retention rates.

4.3 Using Web Analytics Engines for Display Analytics Reporting

Signage analytics can benefit from the wide range of reports and aggregations provided by common Web analytics systems. As part of our exploration into the creation of analytics reports relevant for the digital signage domain without 'reinventing the wheel' and with the purpose to reuse the body of work existent in the Web analytics domain, we have created a possible mapping from signage analytics terminology to Web analytics and developed a corresponding injection module for Pheme. We subsequently used the developed mapping and injection module to create a set of novel display-oriented analytics reports by leveraging on an existing Web analytics engine.

4.3.1 Overview of Web Analytics Terminology

We first provide an overview of the capabilities that state-of-the-art Web analytics engines provide – before developing an appropriate mapping of signage analytics to Web analytics terminology.

Modern Web analytics emerged that implemented comprehensive on-client data collection using JavaScript embedded in Web sites allowing to gain deeper insights into user behaviour and interactions conducted on a single page – such as the Google Analytics tracking library

Table 4.5: Attributes provided to describe Page View and Event hit types in Google Analytics based on the Universal Measurement Protocol (UMP) [Goo18h].

Page View Hit Type	Event Hit Type
Time on site	Category, action, label, value
Bounce rate	Grouping across attributes
Funnels	Non-negative integer values
Entry & exit points	Graphs visualisation for values
Content	Link events to pages
Hierarchical URI / drill down	User ID
Value is 1	Time
Referrer	
Graph visualisation for page views	
User ID	
Time	

[Goo18f]. In the Web, such JavaScript modules are capable of tracking similar user activity to what was achieved through access logs, and beyond by tracking the user's activity within a single page (e.g. reporting button clicks and scrolling behaviour). In addition, powerful analytics engines have emerged that automatically compute relevant metrics and insights upon integrating a simple code snippet on the Web site. An overview of different attributes that can be collected through modern Web analytics as part of page view and the more generic event reporting types are provided in Table 4.5.

In Web analytics, the following key sets of metrics and insights are typically created and offered to Web administrators and content providers.

Sessions A session often describes the thread and lifetime of a set of subsequent user interactions conducted within a certain time frame. In typical Web analytics systems such as Google Analytics the maximum time between the subsequent interactions is defined as 30 minutes [Goo18g] – any interaction of the same user that has been conducted in this time frame is associated to be part of the same session. Sessions are key for the computation of other metrics and reports such as user interaction funnels, i.e. the knowledge of the start and end of a coherent interaction thread is required to create a funnel representing navigation patterns of single sessions.

Pageviews Pageviews are defined as a metric that describes "the total number of pages viewed" [Goo18i] where a single pageview is defined as "an instance of a page being loaded (or reloaded) in a browser" [Goo18i]. Pageviews can be associated with additional attributes such as the time spent on site to provide further insights (Table 4.5). In this case, pageviews are aggregated across pageviews from any user, i.e. if individual users load a particular page multiple times.

Visits Visits count the number of individual visitors that have accessed a website. While pageviews typically aggregate over *any* pageviews that occur, including multiple pageviews from a single user, a *visit* describes the number of sessions a user has started on a Web site [Goo18i]. Visits can be further broken down to distinguish between *unique visitors* (i.e. counting unique visitors only once in a given time frame) and *recurring visitors* (i.e. only considering visitors who have previously visited a Web page).

Bounce Rates Bounce rates define the proportion of visits that are not followed by any subsequent pageview of the same user within the domain, i.e. the proportion of visitors that 'bounce' away from the website after only opening a single page [Goo18b]. Bounce rates are used as a indication of the interactivity of both the Web site and visitors. Depending on the nature of the Web site, low or high bounce rates may indicate a problem or just reflect intended behaviour [Goo18b].

Funnels Funnels (initially described in Section 2.4.2.1, p. 35) typically visualise an interaction or behaviour flow. For example, funnels in Web analytics are used to show an overview of the overall traffic flow within a Web site – i.e. which proportion of visitors views a sequence of Web pages [Goo18f]. Funnels can also be utilised to visualise the flow of custom-defined events or e-Commerce-related purchasing patterns.

Referrers and Traffic Sources Referrers and traffic sources describe the originating or referring Web site of a *visit* or *pageview* event. Examples of referrers and traffic sources can be other, external and internal Web sites as well as direct traffic that has not been referred from another Web site through a automated forwarding or hyperlink [Goo18i].

Conversion Rates Conversion rates describe the proportion of viewers that have completed a *goal* that was previously defined by the Web site administrator or content provider. Examples of goals can include "a completed sign up for your email newsletter (a Goal conversion) or a purchase (a Transaction, sometimes called an E-commerce conversion)." [Goo18a]

4.3.2 Mapping from Signage to Web Analytics Terminology

We developed a mapping that emphasises the similarities between the digital signage and Web analytics domains. In order to report pageviews and events to Web analytics systems, a core set of attributes need to be mapped which typically include meta data about the document that has been opened (including title, domain and URI), the user performing the event (e.g. a globally unique user identifier), and the, if applicable, the referring Web page. In digital signage, we can simply use some of the core metrics recorded as part of signage deployments and map these directly across – an overview of the mapping of individual attributes is shown in Table 4.6. The mapping presented leads to changes in the meaning of metrics reported through Web analytics as shown in Table 4.7. Pageviews and page impressions can be directly mapped

Table 4.6: Mapping of individual attributes from digital signage analytics to Web analytics' attributes (based on UMP [Goo18h]).

Web Analytics Attribute	Digital Signage Analytics Attribute
Tracking identifier	Tracking identifier is passed through
Query time	Time delta between the actual content change (reported by individual displays) and the mapping and injection to the Web analytics platform (allows to account for potential processing and mapping delays)
Client identifier	Unique identifier of the display reporting analytics
Hit type	The hit type that is reported; 'pageview' for content reports or 'event' for more generic analytics reports
Document location	URI to the reported content that is shown on a display
Document host	Source host of the content reported (if available)
Document path	Path to the content reported (if available)
Document title	File name of the content reported
Content description	File name of the content reported

Table 4.7: Mapping of reports from Web analytics to digital signage analytics' metrics (partially described in [Mik+15]).

Web Analytics Metric	Digital Signage Analytics Metric
Page views / page impressions	Content impressions on displays
Referrers	Identifier of display showing the content
Unique visitor	Individual display active as part of the display network and reporting content
Visitors	All displays active as part of the display network and reporting content
Session	Active period of individual displays (i.e. a new session is started if the display has not reported content for a certain time period)
Time on site	Duration of a display showing an individual piece of content
Page title	Title or file name of the content that was shown
Page URI	Source URI (local or Web-based depending on the display) to the content item shown
Location	Geographical location of the display
Browser	Renderer used to show the content
Operating system	Underlying operating system of the display and the signage player software

onto content impressions; the metadata about content shown on displays can be mapped onto the document metadata (including content titles and URIs). Whilst in this scenario we do not have individual users, we use the user metadata attributes to distinguish between individual displays accessing content. Therefore, user identifiers are used to provide identifiers of the display accessing a piece of content. To understand content change patterns and provide us with the ability to filter for individual displays, we were required to make the reporting display visible in the analytics dashboard. Traditionally, in Web analytics systems individual users are anonymised and invisible, and are only reported as part of larger aggregates. In the case of digital signs, such an anonymisation of displays is not required. Instead, we deliberately make display identifiers visible by mapping these values onto the referrer attribute which can be provided for any Web event hit type – causing the display identifier to directly appear in reports and provide us with the ability to conduct further filtering.

We note that the mapping provided in Tables 4.6 and 4.7 is focussed on a display-centric mapping, designed to capture and map across predominantly data from digital signs not considering the audience. However, other mappings depending on the purposes and desired results can be possible. For example, a closer relationship between users of the digital world (i.e. Web users) and physical world (i.e. passers-by) could be directly mapped translating 'display content view impressions' to 'page views'. The results, e.g. number of impressions for given Web sites would directly translate to the number of views of certain content items displayed on digital signs across a signage deployment. In order to support such reports, digital signs need to be able to capture glances and views, e.g. through the use of visual analytics and face recognition solutions. These technologies could be combined with other recognition systems such as Tacita to support more comprehensive metrics including 'unique visitors' of content items – presupposing that the re-recognition of individuals can be achieved across spatially distinct displays and deployments.

4.3.3 Example Reports and Visualisations

Using Pheme, we are able to put analytics into existing digital signage deployments and utilise Web analytics as an example for third-party visualisation and reporting purposes through the mapping described above. We note that the visualisations and aggregations in the following subsections were solely the result of the mapping and did not require additional fine tuning of the dashboard – further emphasising the conceptual similarities between the Web and digital signage analytics domains. The visualisations draw on the data modelling described as part of the display-oriented data collection in Section 3.3.2 (Client-side Data Collection, p. 60).

4.3.3.1 Display-oriented Performance Reports

Figure 4.12 shows the incoming data stream as part of the Google Analytics real-time dashboard. Digital signs reporting content appear as 'active users' and referrers in the dashboard together with the currently shown piece of content in the form of a string. Due to the mapping of display identifiers onto the user identifier attribute, the number of 'active users' directly

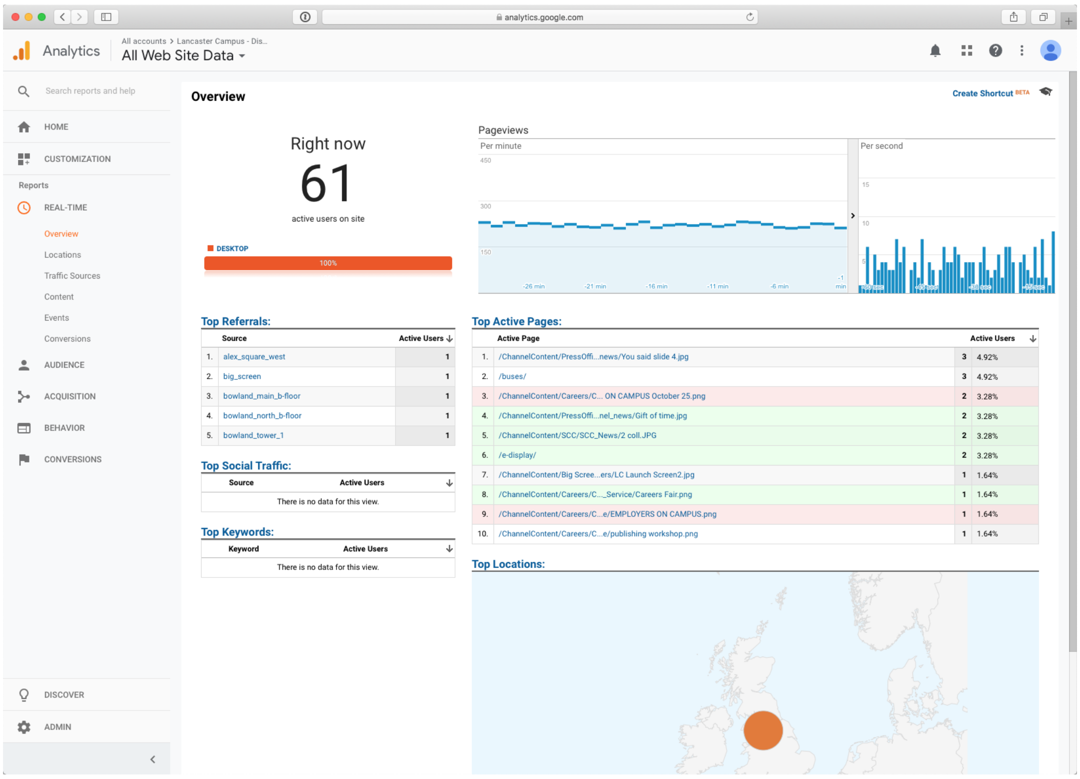

Figure 4.12: Pheme real-time reports produced through Google Analytics.

maps onto the number of currently active displays (i.e. displays that are showing content). Additionally, the list of referrals shows the names and identifiers of active displays. Standard features of the dashboard can be used: clicking on a referrer results in a filtering of real-time insights about the selected display.

As described in Section 3.3.2 (Client-side Data Collection, p. 60), we additionally utilise the custom Events hit type to describe the physical power state of displays to detect malfunctioning displays and signage players. As shown in Figure 4.13, the analytics dashboards provides an overview of reported event types (e.g. 'unresponsive') and any other customisable event type. The analytics dashboard consists of the ability to filter for any custom event type and value and retrieve the set of referrals which directly map onto the corresponding displays. The ability to create aggregates and historical reports can be used to, for example, capture the reliability of a signage network over time. Such insights are crucial for the success of a display network as they allow network administrators and providers to easily determine malfunctioning devices.

4.3.3.2 Specific Reports for Content and Service Providers

In open display network scenarios, content providers do not necessarily know the distribution of their content items across display networks. However, understanding the content display patterns for individual content items across single or multiple display networks is a crucial piece of information for both display owners and content providers – helping to inform the

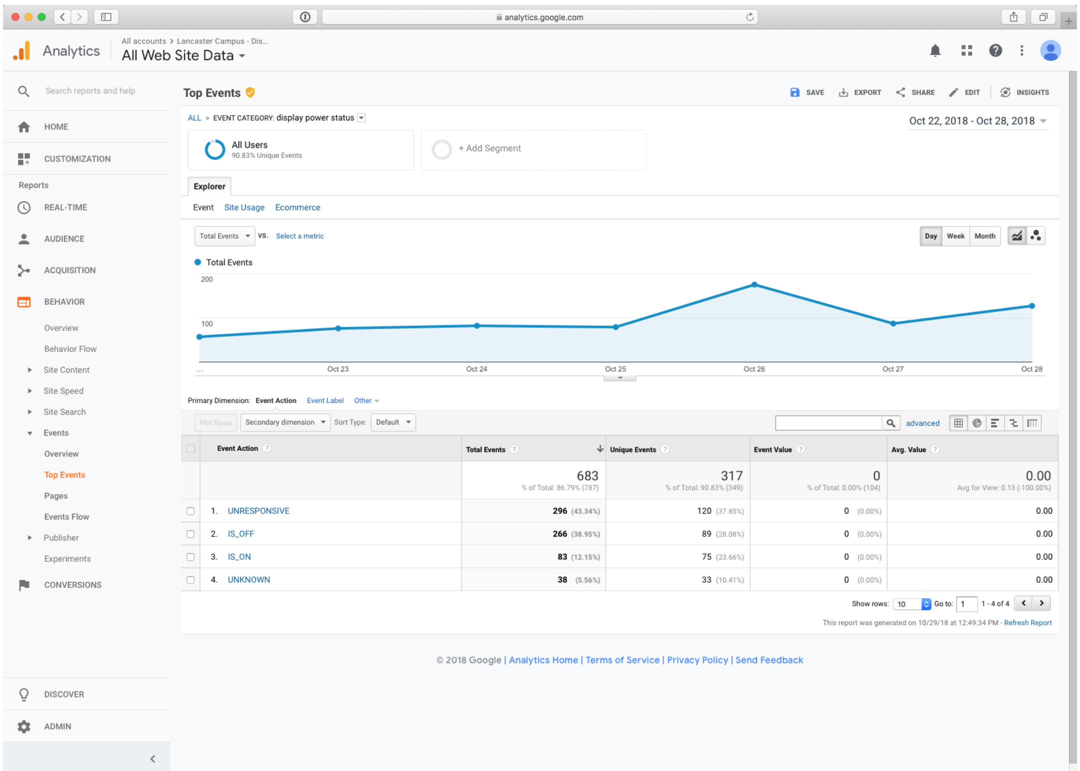

Figure 4.13: Pheme event reports produced through Google Analytics.

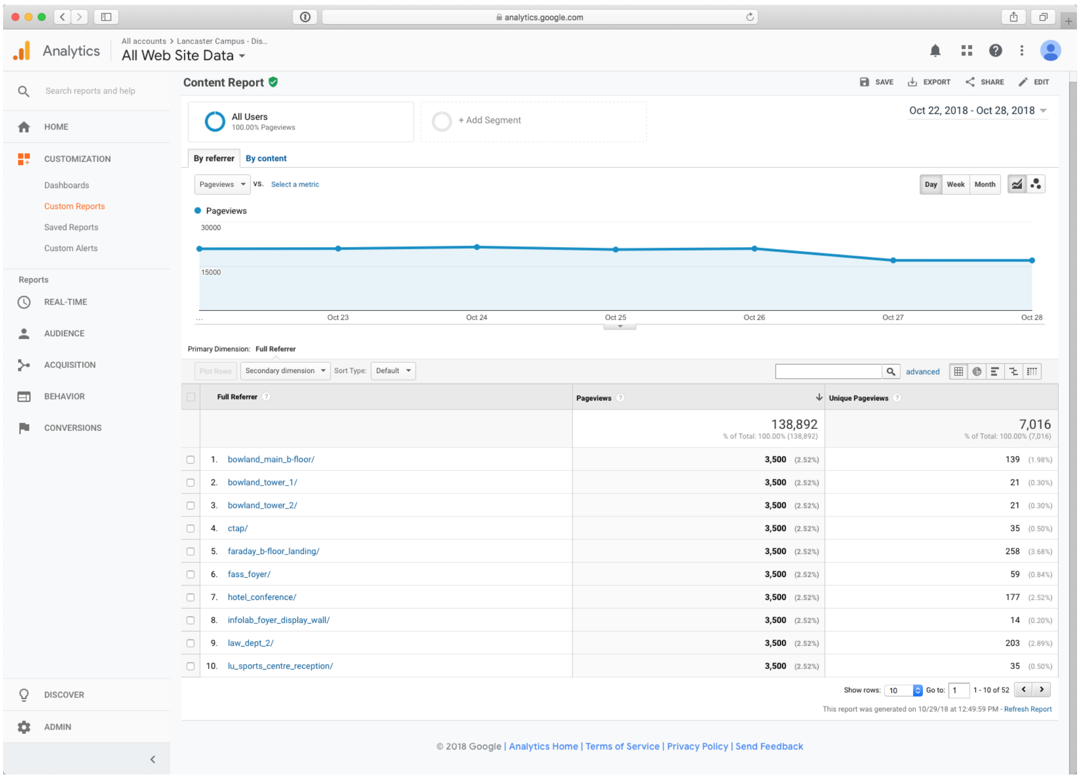

Figure 4.14: Pheme content reports produced through Google Analytics.

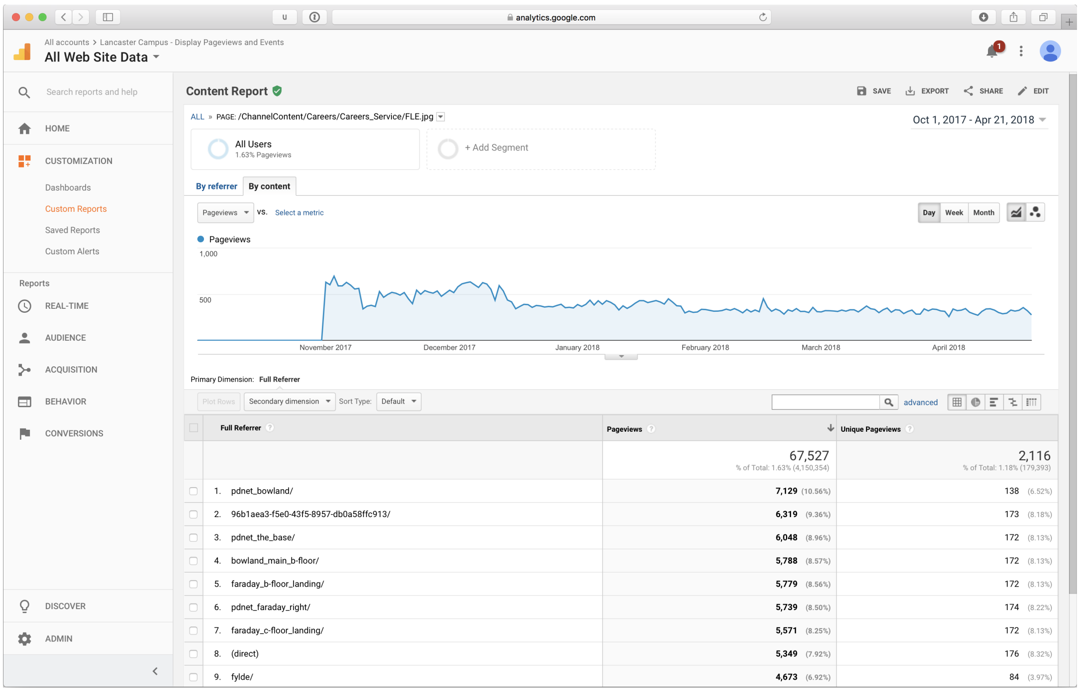

Figure 4.15: Detailed report of displays showing a particular piece of content produced through Google Analytics.

understanding of both the reach of content and the general performance of a signage network. Such reports are particularly important for content providers (as described in Section 2.4.1) to verify that their content has been played across a signage network in line with agreements with display owners. By applying the example of a basic mapping of the previously described pageviews, we are able to produce comprehensive content reports and aggregations throughout the entire lifespan of a display and piece of content. Figure 4.14 shows the overview of displays (i.e. referrers) that were reported showing content, allowing administrators to drill down on a per-display basis or filter for specific content items instead to identify which display has been showing a particular piece of content at which times. For example, Figure 4.15 shows an overview of the distribution and impressions of a particular piece of content across the entire digital signage network – allowing both display owners and content providers to understand the visibility scope of particular pieces of content.

4.4 Summary

In this chapter, we introduced a set of novel viewer-centric analytics reports for digital signage based on our approach of combining display-oriented analytics with viewer mobility data. Concretely, we have made the following contributions:

1. We have presented a set of novel viewer-centric analytics reports that illustrate the levels of insights that can be gained when considering viewer sightings of displays across a network of displays – including reports showing the effectiveness of displays, the

visibility of content across the display network, and the visibility of content to viewers. Our analytics reports were drawing on both synthetic mobility traces and the display sightings captured through Tacita.

2. We used Tacita as a use case to presented a set of analytics reports specific to supporting a display personalisation system. In particular, such reports included reports regarding usages and interactions, and an overview of using ‘retention rates’ (and its limitations).

3. We illustrated that leveraging existing Web analytics engines enabled us to create a set of novel sign-oriented analytics reports. In particular, we presented a mapping from signage to Web analytics terminology, and designed and developed an appropriate injection module for Pheme that allowed us to report display and content related reports for the e-Campus display network.

In many analytics systems the end product is the set of reports provided. However, analytics data can also be used to inform the content shown and the behaviour of the signage network. In the following chapter, we describe the automated use of collected and computed analytics data on pervasive displays, e.g. to improve the quality of the network and the viewer experience.

Chapter 5

Automated Use of Pervasive Display Analytics

5.1 Overview

In the previous chapters we focussed on data capture (Chapter 3) and reporting aspects (Chapter 4) of next generation display analytics systems. However, analytics insights could also be used to drive content scheduling decisions on digital signs. In this chapter, we focus on the design and development of a novel content scheduling system and present the Lottery Scheduler, a new approach to dynamic content scheduling that supports both traditional content scheduling and provides the ability for context- and event-based scheduling. In particular, we describe how the Lottery Scheduler can be used as a solution given the high number of potentially conflicting content scheduling constraints and requirements that are likely to emerge in future open pervasive display networks that use analytics data to inform the content selection.

Excerpts of this chapter have been published in the following peer-reviewed publication:

1. Mateusz Mikusz, Sarah Clinch, and Nigel Davies. "Are You Feeling Lucky?: Lottery-based Scheduling for Public Displays". In: *Proceedings of the 4th International Symposium on Pervasive Displays*. PerDis '15. Saarbruecken, Germany: ACM, 2015, pp. 123–129. ISBN: 978-1-4503-3608-6. DOI: 10.1145/2757710.2757721. URL: http://doi.acm.org/10.1145/2757710.2757721

5.2 The Need for Dynamic Content Scheduling

A key challenge in a future analytics-driven signage system is how to incorporate the analytics insights into content scheduling decisions. In this section, we introduce existing content scheduling approaches from the digital signage domain, and also provide an overview of basic scheduling algorithms that have been developed in the context of fundamental task and process scheduling research.

5.2.1 Existing Content Scheduling Approaches

Traditionally, content scheduling for public displays is carried out using predefined timelines or playlists that consist of a detailed description of when and where a certain piece of content should be shown – either for individual displays or for a group of displays. Examples of a commercial system with such complex timeline-based scheduling capabilities include Sony Ziris [Son] and BroadSign [Bro] that provided users with comprehensive user interfaces, allowing detailed control of content shown on individual displays within the signage network. The use of a combination of scheduling techniques has been used in a number of research works including [Fin+96; MCL01; Chu+03; Elh+14]. For example, researchers often utilised interaction-driven content scheduling (i.e. based on explicit input of the viewer through mobile phones [Dav+09]) combined with simple cycling through a predefined set of content items [KGR08]. Further research, such as work conducted by Storz et al. [Sto+06] and Elhart et al. [Elh+14], explored the design and development of scheduling systems that allow the specification of a range of scheduling constraints including a server-based scheduler that manages the content selection on a network of displays. Elhart et al. [Elh+14] further developed a scheduling language that allows the specification and formalisation of complex content scheduling constraints and requirements in digital signage.

The scheduling approaches above typically lack consideration for contextual and other events that may influence the scheduling decision of current or future content. The need for digital signage players that support dynamic interventions and content changes becomes crucial when considering contextual information alongside analytical insights for content scheduling decisions. For example, digital signage systems can learn and adapt to a certain behaviour or an audience that is currently moving through a space by passing on analytics information of the audience to digital signs situated in their proximity. A digital signage player deployed in the space could consider such real-time insights and use that information to dynamically inform the scheduling decision of the currently shown content item. Whilst state-of-the-art signage players only support intervention for a very limited set of context such as direct content interactions and input, the support for the described scenario would require a highly dynamic and flexible scheduling system.

5.2.2 Scheduling in Operating Systems

The problem of scheduling content onto a public display can be seen as a resource allocation problem: a number of content items (likely to have originated from a distinct set of stakeholders) representing ‘tasks’ compete over limited screen real-estate representing the ‘resource’. In the context of operating systems, a number of resource allocation and process scheduling techniques have been developed considering varying constraints and requirements – some of these resource allocation techniques have formed the basis for the content scheduling algorithms described above. To provide an overview of fundamental task scheduling techniques, Panwalkar and Iskander [PI77] conducted a survey of techniques developed in early research. The authors of [PI77] categorised scheduling approaches into three overarching classes: priority-based scheduling (tasks are allocated based on specific attributes such as due dates by prioritising tasks with the earliest deadlines first or the number of resources requested), heuristics (tasks are allocated based on more sophisticated mathematical rules that take additional factors into considerations such as intended task loads), and other rules (tasks are allocated based on rules designed for a specific purpose, or a combination of priority-based approaches and heuristics) [PI77]. Specific to operating systems, Arpaci-Dusseau and Arpaci-Dusseau provide an overview of existing scheduling approaches including ‘first in, first out’, ‘shortest job first’, ‘shortest time to completion’ and ‘round robin’ (switches between jobs after each process execution cycle, i.e. tasks receive an equal amount of time slices to complete their jobs) [AA15].

Examples of more sophisticated scheduling algorithms include *Stride Scheduling*, a “deterministic allocation mechanism for time-shared resources” [WW95] developed by Waldspurger and Weihl. W. In this approach, resources (over which a number of tasks are competing) are allocated ‘deterministic time slices’ whilst the access to such resources is represented by tickets. Tasks that have been allocated a number of tickets hold ‘access rights’ to these resources. In return, the scheduler executes (competing) tasks in strides that are inversely proportional to the number of tickets a task holds; i.e. a task with twice as many tickets is given twice as much access to resources compared to other tasks [WW95].

An alternative approach for the distribution of tasks in a resource-constrained context is *Lottery Scheduling*, first introduced as a“flexible proportional-share resource management” [WW94] approach by Waldspurger and Weihl. The system was motivated by the challenge to schedule a large set of computational tasks competing for a limited set of available computing resources. To provide a way of ‘fair’ scheduling and distribution of tasks, the lottery scheduling approach provides a mechanism for modelling rights to resources in the form of lottery tickets that are allocated to tasks waiting for resources. Once the allocation of lottery tickets is complete, a lottery is held to determine the allocation of resources to tasks. Waldspurger and Weihl state that this approach “effectively allocates resources to competing clients in proportion to the number of tickets that they hold” [WW94], i.e. providing a fair yet effective way to distribute resources to tasks reflecting their requirements. Changes to the allocation of lottery tickets reflecting changes in the resource requirements are immediately reflected in the subsequent draw allowing the Lottery Scheduler to dynamically and quickly respond to changes.

In addition to the proportional distribution of lottery tickets, Waldspurger and Weihl developed additional, modular mechanisms to allow tasks that influence the resulting draw [WW94]. First, the scheduler supports the transfer of lottery tickets between tasks. For example, if a task is waiting due to another task using up or blocking resources, the pending task may temporarily transfer tickets to the blocking task to enable a faster execution. In order to enable high prioritisation of certain tasks, the authors provide the ability for a ‘client’ (i.e. the entity holding resources) to create a large number of tickets and allocate these to the task that is to be prioritised—defined as “ticket inflation”. For more detailed resource allocation within a single client or task, the scheduler allows the creation of internal “ticket currencies” used within an internal group of tasks. Finally, Waldspurger and Weihl provide "compensation tickets" to clients that utilise only a subset of allocated resources to guarantee a proportional distribution of lottery tickets and proportions in future lottery draws [WW94].

5.3 Lottery Scheduling for Digital Signage

In order to facilitate the requirements for a dynamic scheduling system that can both consider scheduling constraints defined by users as well as react to frequent changes in the context of the sign, we have designed a lottery-based scheduling system for digital signs that builds on top of the approach initially introduced by Waldspurger and Weihl [WW94]. The lottery scheduler is highly modular and configurable and has the ability to resolve complex scheduling requirements and constraints though the use of lottery-based algorithms.

5.3.1 Applicability of Lottery Scheduling to the Public Display Domain

The allocation of resources to clients to perform certain computational tasks by an operating system can be transferred into the digital signage domain: the time and screen real estate for content play-out on digital signs is highly limited and typically a resource for which content providers and advertisers compete. The concept of lottery scheduler provides a simple yet effective approach to select a 'winning' tasks whilst allowing for consideration of potentially conflicting scheduling constraints and requirements in the form of lottery tickets and ticket allocation techniques. The lottery scheduler can be easily adapted into the digital signage domain: competing stakeholder requirements can be represented by varying lottery ticket allocation techniques that allocate lottery tickets to available content items. By performing a random draw on the available set of lottery tickets, the digital sign can find at any given point in time an eligible content item to show – whilst the distribution of lottery tickets shifts the probabilities for certain content items to be shown with respect to the requirements and constraints provided by stakeholders. In comparison to other scheduling approaches such as stride it provides a simple solution yet enough flexibility to consider additional scheduling constraints and requirements. We believe that using the lottery scheduling approach will allow the break-down of complex scheduling decisions into smaller, distinct problems formulated as a collection of lottery ticket allocators that contribute to the overall scheduling decision. Additionally, a number of external and contextual events may influence content scheduling decisions at any given point of time, introducing an additional level of complexity regarding the scheduling decision and mirroring the need for ticket re-allocation seen in operating systems. We emphasise that, however, the lottery scheduling approach solves the issue of determining which content item to show from a set of *eligible* content items. Common content scheduling systems for public displays allow the definition of additional scheduling constraints, e.g. making certain content items available at certain days or time ranges only [SFD06]. Whilst such scheduling constraints could be implemented by allocating zero tickets to ineligible content items, we believe that such an approach would create an overhead as each individual lottery ticket allocation component would be required to identify the eligibility of individual content items in addition to determining the number of lottery tickets to be allocated. In order to illustrate the applicability of the lottery scheduling approach for public displays, we describe how the Lottery Scheduler can be used to address a set of typical scheduling requirements that emerged based on experiences from the e-Campus display deployment [SFD06; Sto+06].

Ratio-based Scheduling Ratio-based scheduling describes the ability of display owners to express ratio-based preferences as to which content receives proportionally more screen time compared to other content. For example, display owners may prefer to prioritise content supplied by themselves over content that is supplied by the space owner or a separate organisation. In the context of the Lottery Scheduler approach, ratios can be simply expressed through the proportional allocation of lottery tickets to certain content items. Of course, factors such as the length of a content item (i.e. for videos) needs to be taken into consideration when allocating an appropriate amount of lottery tickets to calculate the screen time correctly.

Viewer Linger Times Viewer linger times serve as an example for considering contextual events taking place in the immediate vicinity of individual displays. In some cases, viewers may linger for longer or shorter periods of time in front of a display which may be considered for content scheduling purposes. For example, displays located near a shop may attract only short glances, whilst displays located in waiting areas at airports or train stations have the potential to longer glances of passers-by. The use of lottery scheduling allows the system to dynamically adapt to viewer linger times and schedule content accordingly to, for example, increase probabilities that certain content items are seen by the passers-by.

Prioritisation of New Content In digital signs with a large volume of content could use a ticket allocation strategy based on the age of individual pieces of content to prioritise more recent content. The advantage of such an approach minimises the delay in which more recent content appears on digital signs and make it visible to viewers. The prioritisation of new content can be implemented by simply allocating proportionally more lottery tickets to newer content, e.g. based on a linear or exponential function depending on the desired output.

Targeted Content and Personalisation Display personalisation systems, such as Tacita [Dav+14; Mik+18d] provide display owners and content providers with the ability to dynamically tailor the content to the passers-by and additionally support long-term personalisation. For dynamic personalisation, multiple passers-by are likely to be competing over screen time with each other and the ‘regular’ content schedule for a display. The use of a lottery ticket allocation based on requested items and a random draw allows the sign to quickly determine a piece of content to show whilst still taking competing content requests under consideration. For supporting long-term personalisation (i.e. considering viewer presence in a space or area and adapting the

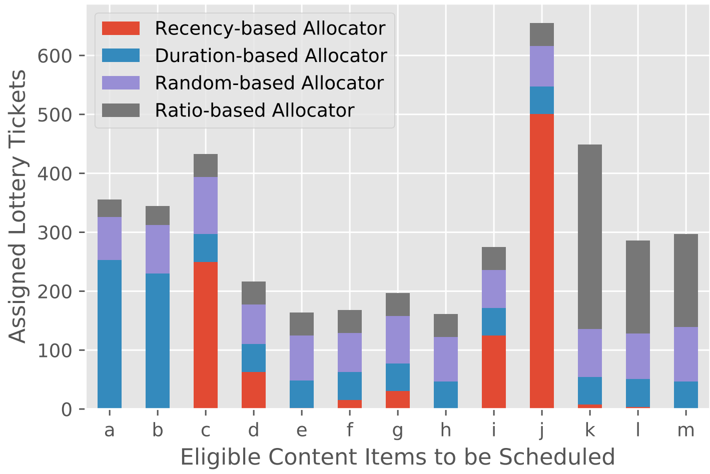

Figure 5.1: Snapshot of an allocation of lottery tickets for 13 distinct content items using recency-, duration-, random- and ratio-based lottery ticket allocation modules each representing different stakeholder requirements.

content over a longer time [Dav+14]), a similar approach can be taken in which lottery tickets are allocated based on the presence of viewers in the wider context of a display.

The lottery scheduling approach further supports the implementation of complex scheduling scheduling requirements by mixing a set of lottery ticket allocations strategies. For example, to support both long-term personalisation and ratio-based scheduling, a mix of multiple lottery ticket allocation modules could be applied at the same time. In this case, a single scheduling system would support both long-term scheduling requirements and these that require immediate changes to the scheduling approach. Of course, the use of different lottery ticket allocation approaches influences the overall probability of certain content items to be scheduled onto the display and needs to be further considered when applying a mix of lottery ticket allocation approaches.

Figure 5.1 provides a snapshot of an example lottery ticket distribution for a set of eligible content items using four lottery ticket allocation mechanisms in order to help to express how the lottery scheduling approach can be utilised to resolve potentially conflicting scheduling requirements from distinct stakeholders. In this example, lottery tickets are allocated using the following strategies: recency-based (i.e. lottery tickets are inversely proportional to the last time the content was played), duration-based (i.e. shorter content items are prioritised by assigning lottery tickets in proportion to the duration of the content), random (i.e. lottery tickets are randomly allocated), and ratio-based (i.e. with respect to predefined content ratios to prioritise certain content over others). Each lottery ticket allocation has been given a total of $1,000$ lottery tickets that are then allocated based on the methods described. Each lottery ticket allocation mechanism can be a direct representation of different stakeholder requirements. For example, the recency-based lottery ticket allocation can represent the requirement of content providers to prioritise distribute play times of content evenly and therefore prioritise content that has not been shown. In contrast, display owners may prefer that content is played in accordance to a predefined set of ratios to prioritise local content over other content. As a result, a number of requirements likely conflicting can be defined in a set of individual lottery ticket allocations to content items – directly impacting the probability of individual content to be selected. The random draw enables a prompt selection of content to be shown despite the set of conflicting scheduling requirements. For example, in the context of personalisation, despite user requests, other content can be still scheduled on the display ensuring that viewers cannot overtake displays for a long period of time – no disadvantage for users who have not opted in for display personalisation. We note that Figure 5.1 shows a snapshot of lottery ticket allocations at a specific point in time – lottery tickets are constantly reallocated in order to account for potential changes in the context that may be considered by a lottery ticket allocator.

We note that the lottery scheduling approach in the context of digital signage also imposes a set of limitations as certain content scheduling requirements cannot be met. For example, the visualisation of content items in a predefined order cannot be guaranteed when a drawing is used to determine the next item to show – even if the probabilities have been adjusted by allocating a corresponding amount of lottery tickets. More generally, if display owners require to show content items as part of a timeline in which full control is given for exact dates and times in which certain pieces of content are played, their specific ordering, and potentially other demands, the lottery scheduling approach would not be an appropriate scheduling solution.

5.4 Lottery Scheduling System Architecture

5.4.1 System Architecture Overview

Considering the requirements above, we designed and developed the first Lottery Scheduler for the digital signage domain. This consists of a series of components as shown in Figure 5.2. As described in the previous section, content management systems typically support the specification of basic constraints defining the availability of individual content items (e.g. based on date and time). In order to simplify the design of individual lottery ticket allocation modules, we integrated an optimisation into the system architecture of the Lottery Scheduler in which the filtering component removes ineligible content items before the execution of the lottery ticket allocation components. By applying this optimisation, we reduce the level of complexity and duplication across other components – lottery ticket allocation modules are not required to additionally review the basic eligibility of individual content items. The Lottery Scheduler executes the following set of processes in sequence:

1. constraints processing on the given set of scheduled content items and filtering of content items that are not eligible to show;

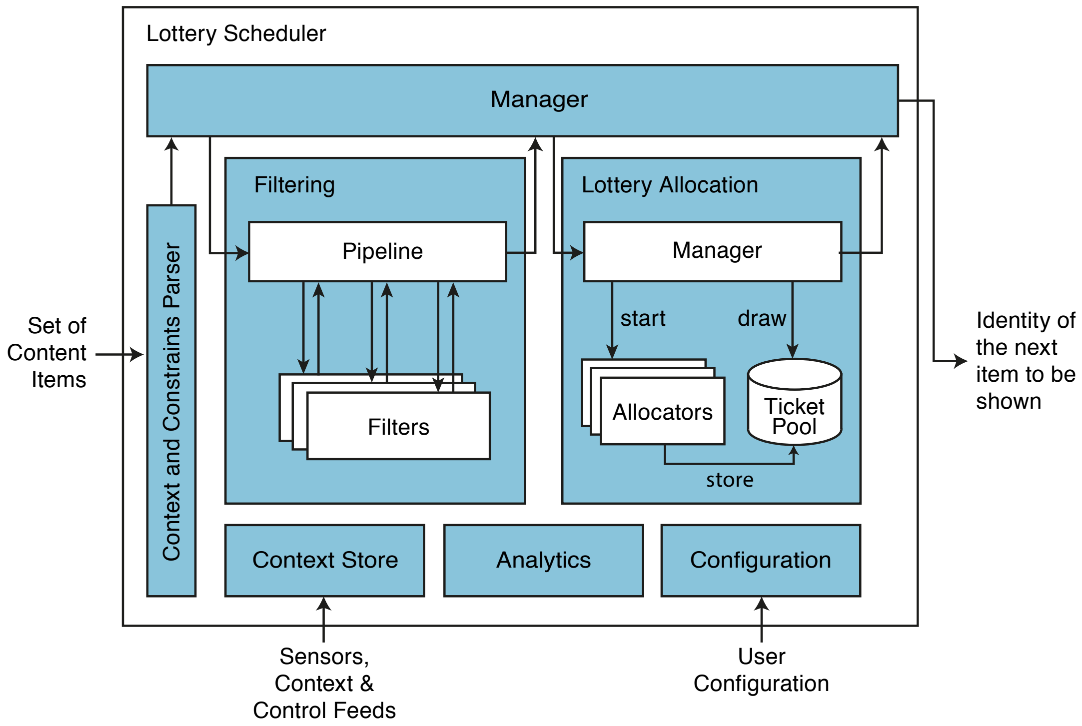

Figure 5.2: Lottery Scheduler system architecture (initially published in [MCD15]).

2. lottery ticket allocation to the remaining set of content items based on strategies defined within each lottery ticket allocator, and

3. random lottery draw; the content item associated with the winning ticket becomes immediately visible on the display.

The Lottery Scheduler consists of the following components: *scheduling pipeline manager*, *context and constraints parser*, *filtering* and *lottery ticket allocation* pipelines, *configuration*, *context store* and an *analytics* module. Each of these components are described in more detail in the subsequent sections.

5.4.2 Scheduling Manager

The Scheduling Manager component orchestrates the data flow across all system components and consists of event-based interfaces for receiving and processing updates from other components such as the context and constraints parser and sensor manager if the content schedule description was updated or if relevant contextual events from online or physical sensors have been reporter respectively. The manager component initiates new lottery scheduling draws when necessary and interacts with the remainder of the digital signage player ecosystem to pass on 'winning' content items to be shown on the display.

5.4.3 Context and Constraints Parser

The Context and Constraints Parser is responsible for pulling and processing content schedules for the individual display. Such content schedules are typically retrieved from web-based interfaces, for example, in the form of an XML-based Content Descriptor Set accessed through the Channel system (initially described in Section 1.3, p. 1.3). Retrieved content schedules are transformed into a standardised, object-based format that can be processed from the subsequent components of the scheduling system. The Context and Constraints Parser is highly modular allowing additional parser modules to be written to interact with APIs of other back-end systems that use a different format to describe content schedule descriptions.

5.4.4 Context Store

The Context Store acts as a repository for analytics data and contextual events captured through internal or external sensors. The ability to react to such data to provide dynamic scheduling decisions is one of the core advantages of the lottery scheduling system. The definition of a 'sensor' in this case is broad and include external Web-based APIs and physical sensors connected to the sign. Any component part of the Lottery Scheduler ecosystem can feed information into the context store. For example, the context and constraints parser or sensors may provide real-time information about current audience counts present in the vicinity of the display. Other components of the Lottery Scheduler such as individual filters and lottery ticket allocators can access any information stored in the context store and utilise the information to inform their filtering and ticket allocation decisions. Both filter and ticket allocation modules can feed information into the context store themselves such as recent allocation decisions and winners of the random draw to support more 'intelligent' lottery ticket allocation based on previous decisions (e.g. to support the prioritisation of most recently added content items).

5.4.5 Filtering

The Filtering component identifies the set of *eligible* content items from the total set of content items scheduled for the particular display. Within the Filter, data flows through a predefined set of filters. In order to support a wide range of constraints, we have designed the Filtering component following a pipeline model in which a series of individual filters are called in turn. Each individual filter returns a set of eligible content items; and upon passing through all available filters, the output of the filter pipeline is the intersection across all sets returned from each Filtering component. Any of the remaining content items are eligible to be shown on the display at the given point in time and contextual state. The resulting set of content items is returned back to the Scheduler Manager where it will subsequently be passed to the Lottery Scheduler component.

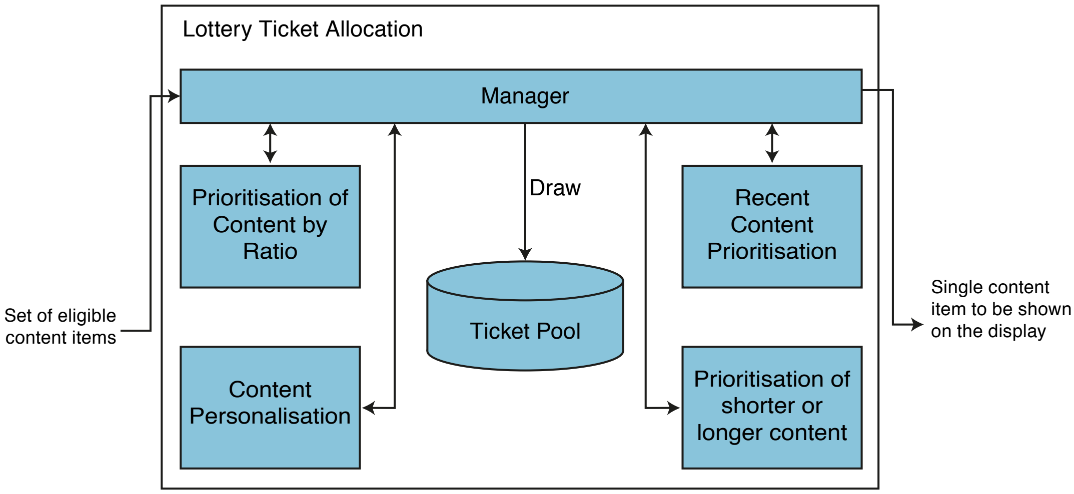

Figure 5.3: Ticket allocation component of the Lottery Scheduler (initially published in [MCD15]).

5.4.6 Lottery Ticket Allocation

The lottery ticket allocation component is the core component of the overall lottery scheduling ecosystem and responsible for both the allocation of lottery tickets and the random draw at the end of the process. This is the next component in the pipeline after the filtering and receives a set of content items, any of which eligible to be shown on the display at the point in time. To represent different, co-existing scheduling requirements, the Lottery Scheduler provides the ability to include multiple lottery ticket allocation modules which can run simultaneously (Figure 5.3).

More specifically, the lottery ticket allocation component maintains a common *lottery pool* into which each lottery ticket allocator can add their lottery tickets. Each lottery ticket is associated to a single content item from the set of eligible content items. Each lottery ticket allocator operates autonomously but is allocated a fixed set of lottery tickets to distribute and access to the set of eligible content items by the lottery manager. In addition, each lottery ticket allocator has access to the context store and configuration component. Whilst each lottery ticket allocator could take an arbitrary amount of time to complete the ticket allocation process, the manager has the ability to interrupt individual lottery ticket allocators after a given time threshold and perform the draw on the allocated set of lottery tickets. To avoid an unwanted interruption, ticket allocators can indicate their internal state and report if the ticket allocation process has finished. Once the scheduling manager has decided that the draw will take place (either due to exceeding the maximum time for ticket allocations or if all ticket allocators have reported a ‘ready’ state), the manager will perform a random draw from the available pool of tickets. The content item associated with the winning ticket will be returned to the manager class and immediately shown on the display.

The user of the lottery scheduling system (typically represented by the display owner) has the ability to decide which lottery ticket allocation modules are used. In addition, the system allows the specification of the number of lottery tickets each module can access to allocate and add to the pool of lottery tickets. This provides administrators with the ability to influence the scheduling decision and, for example, prioritise certain scheduling decisions over others by allowing ticket allocators to only access a limited amount of lottery tickets. Due to the modular design, new lottery ticket allocators that implement new scheduling strategies can be integrated easily.

5.4.7 Configuration

The Lottery Scheduler has been designed with modularity and configurability in mind allowing display owners and other stakeholders to individualise and influence the behaviour of the overall ecosystem. Preferences regarding the activated set of filter and ticket allocation modules including the specification of a maximum of available lottery tickets are stored within the configuration component and made accessible to any other system components. In addition, default content lengths, frequency of content changes and potential analytics back-ends for reporting content changes can be configured and stored within this component.

5.4.8 Analytics

The analytics component within the Lottery Scheduler is conceptionally designed to capture the state of the Lottery Scheduler and capture both content scheduling decisions and contextual changes and events. This component supports the integration of third-party or external analytics engines such as Pheme (introduced in Section 3.3, p. 59) and reports events in real time.

5.5 Implementation

The Lottery Scheduler has been implemented as a component within Yarely, a digital signage player initially developed by Clinch et al. [Cli+13] and deployed in the context of the e-Campus display test-bed described in Section 1.3 (Research Context, p. 6). Yarely is a highly modular system and consists of a five components relevant for digital signage (Figure 5.4): *Scheduler*, *Subscription Manager*, *Analytics Manager*, *Sensor Manager*, and the *Lifecycle Manager*. All Yarely components communicate through the *Internal Eventing System* implemented on top of ZeroMQ[1], a distributed messaging system allowing isolated components and processes to communicate over low-level sockets. Yarely orchestrates both the retrieval of content schedules for individual displays through the Subscription Manager using the Content Descriptor Set format as well as the playback of content on the actual display.

The Lottery Scheduler was implemented in Python (1,051 lines of code) as a direct replacement of the Scheduler component within Yarely. Each of the Lottery Scheduler components have been implemented as individual Python modules. Within the filter and Lottery Scheduler components, each filter and lottery ticket allocator are implemented as separate Python-based classes with common input and output specifications, enabling an easy

[1] https://zeromq.org

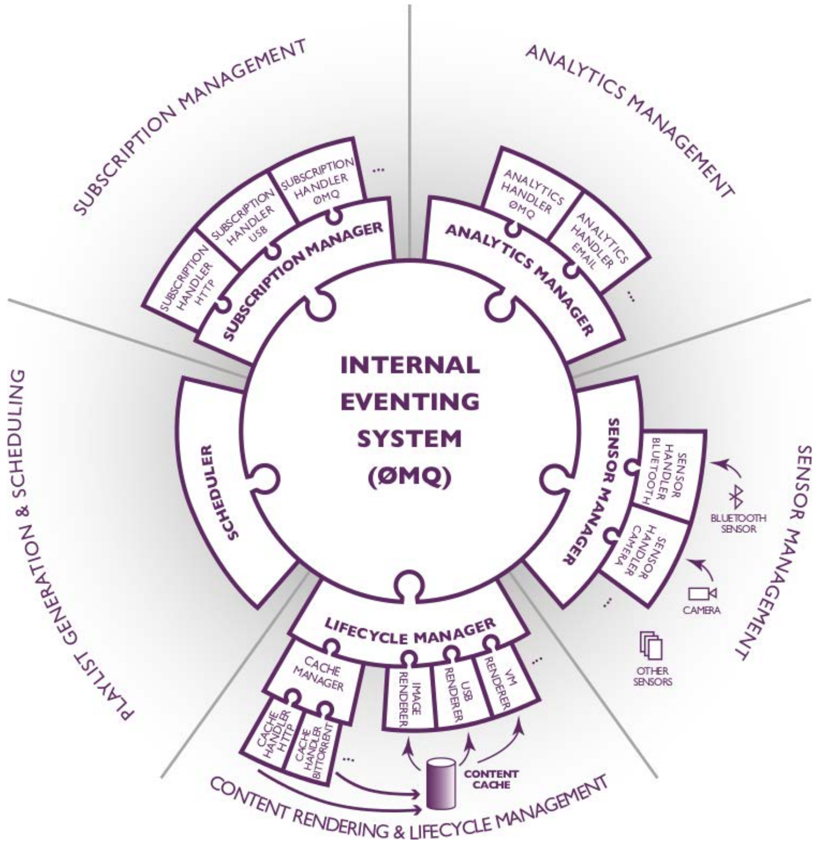

Figure 5.4: Yarely system architecture developed by [Cli+13] (reproduced from [Cli+13]).

extension with new filters and ticket allocators. Whilst filters were implemented to run in sequence, the lottery ticket allocation has been parallelised in that each ticket allocator runs in a separate thread to improve the overall system performance. Lottery ticket associated to individual content items are stored in a common queue from which the manager performs a random draw. The description of the winning content item is passed on through the appropriate ZeroMQ channel to the display component of Yarely in order to make the actual content immediately visible on the display.

5.6 Evaluation

We note that the technical feasibility of using the Lottery Scheduler, its system performance and the integration into an existing public display network is evaluated in Section 6.6.

5.7 Summary

In this chapter, we introduced opportunities for utilising insights and reports gained through signage analytics tools for actuation on the sign. In particular, this included:

1. The motivation for the need of novel scheduling on digital signs in the context of open pervasive display networks in which a high number of potentially conflicting

requirements and constraints influence content scheduling decisions on digital signs, and

2. The design and development of a novel lottery scheduling approach for digital signage providing means to support context- and event-based scheduling to allow feeding analytics back into the sign.

In the following chapter, we will describe small- and large-scale trails evaluating previously introduced concepts and systems, e.g. the Lottery Scheduler system and Tacita.

Chapter 6

Trials

6.1 Overview

In the previous chapters, we introduced systems and concepts that enable us to collect relevant analytics data, create novel analytics reports for the digital signage domain, and benefit from such analytics report by designing underlying technology for systems that are capable of automatically using the collected datasets to improve the effectiveness of signs (e.g. by providing personalised or targeted content to the passer-by).

In this chapter, we focus on the evaluation of the systems introduced in the previous chapters including the data collection, reporting and actuation strands. Beginning with data collection, we provide evidence for the performance and reliability of Pheme that has been integrated in the context of the e-Campus display test-bed and used to inject analytics data into existing third-party analytics engines for report creation purposes. As one of our major contributions of collecting viewer-related mobility data we describe the design of our long-term in-the-wild-trial of Tacita. We provide detailed insights into the accuracy and reliability of using Bluetooth Low Energy beacons for the capture of viewer behaviour in the proximity of digital signs, the system latencies of our backend implementation, and conduct an analysis of application usage patterns. Subsequently, we conduct an evaluation of Wi-Fi-based collection of viewer mobility data focussing on providing insights into the accuracy and reliability of using Wi-Fi fingerprinting for the detection of viewer proximity to displays. We provide a detailed description of a controlled walk-by experiment, reporting on the accuracy and performance of the approach for different types of display location and varying system parameters. In addition, we highlight the benefits and costs of both approaches and discuss their applicability for each stakeholder group. We conclude this chapter by providing an evaluation of the Lottery Scheduling approach and provide insights into its performance and scheduling accuracy in the context of the e-Campus display deployment.

6.2 Pheme: Display-oriented Data Collection

The Pheme architecture has been introduced and discussed in Section 3.3 (Capturing Traditional Signage Analytics Data, p. 59) and Section 4.3 (Using Web Analytics Engines for Display Analytics Reporting, p. 102). In the subsequent sections, we will focus on the evaluation of Pheme and associated injectors by providing a description of the integration of Pheme into a large display network, and evidence for the successful implementation and usage of an injector module that inserts incoming datasets in real-time into Google Analytics, a third-party analytics engine.

Excerpts of this section are based on the following publication:

- Mateusz Mikusz, Sarah Clinch, Rachel Jones, Michael Harding, Christopher Winstanley, and Nigel Davies. "Repurposing Web Analytics to Support the IoT". in: *Computer* 48.9 (Sept. 2015), pp. 42–49. ISSN: 0018-9162. DOI: 10.1109/MC.2015.260. URL: http://doi.org/10.1109/MC.2015.260

6.2.1 Integration into e-Campus

We integrated Pheme into our existing e-Campus test-bed (introduced in Section 1.3, Research Context) allowing us to conduct a long-term evaluation of the collection of an incoming stream of analytics reports and the mapping of requests into third-party analytics services.

Pheme was integrated into the e-Campus display test-bed by utilising the dedicated Pheme client libraries (initially described in Section 3.3.2.2) into the Yarely signage player as a direct replacement of Yarely's *Analytics Manager* module, providing access to all Yarely environmental variables and the context store and enabling the monitoring of content changes and incoming events from sensors and other internal and external sources.

The use of our client library allowed us to integrate Pheme into the existing code base of the signage network with minimal effort. Listing 3 shows a simplified version of the Yarely scheduling procedure with the integration of Pheme: the integration of Pheme consisted of the creation of a new method that allowed the specification of tracking identifiers and configuration parameters (`_initialise_analytics(self)` – in our case this method reads the required configuration parameters from a configuration file) and the call of the `track_pageview_async(self, new_item)` to report a new content scheduling event to Pheme. In order to minimise the latencies between the point at which the content appears on the display and is reported to Pheme, we placed `track_pageview_async(self, new_item)` immediately after the function call that initiates the content to be made visible on the display. Additionally, the `track_pageview_async(self, new_item)` method is executed asynchronously allowing the remainder of the scheduler process to continue without blocking on the HTTP request that is performed to the Pheme back-end in order to report the scheduled content item event. The Pheme client library additionally includes a local timestamp of the event occurrence to account for any potential transmission latencies. In addition to reporting content changes, we integrated `track_event_async(self, category, event, value, label)`

```
from phemelibrary import PhemeAnalytics


class SchedulingManager(ApplicationWithConfig):
""" The Scheduling Manager is the global class controlling the scheduling
of Yarely, including the Lottery Scheduler. It is requesting new content
items from the scheduling component and makes these visible on the display.
"""

    def _initialise_analytics(self):
        analytics_tracking_id = self.config.get(
            'Analytics', 'tracking_id', fallback=None
        )
        self.analytics = PhemeAnalytics(analytics_tracking_id)


    def item_scheduling(self):
      # Perform Filtering
      self.filtered_cds = self.filter_pipeline.filter_cds(self.cds)

      # Initiate Lottery Scheduling
      new_item = self.scheduler_pipeline.get_items_to_schedule(1)

      # Make new content visible on display
      self.display_manager.display_item(new_item)

      # Report content change to Pheme
      # Initialise analytics component if not yet activated.
      if not self.analytics:
        self._initialise_analytics()
      self.analytics.track_pageview_async(new_item)
```

Listing 3: Code snippet of the Pheme display client integration into Yarely (simplified).

into Yarely's *Sensor* module to report dynamic content requests issued as part of the Tacita personalisation system.

To allow more detailed insights into engagement and interaction patterns, we further integrated Pheme into individual display applications to support the tracking of button clicks on touch-enabled displays. The integration is similar to common Web Analytics and supports the on-site tracking of viewer interactions. We note that, however, the e-Campus display network consists of non-touch enabled displays only and the tracking of interaction on touch-enabled displays was only deployed for a single, short-term test.

6.2.2 Mapping and Injection Module Integration

In order to provide evidence into the extendability of Pheme as a platform to inject analytics data into third-party services, we developed an injection component that implements the introduced mapping of sign analytics to Web analytics terminology (Section 4.3, p. 102) – used to inject the stream of incoming analytics data to Google Analytics. In addition to demonstrating the extendability of Pheme and usability of the overall framework, the injector

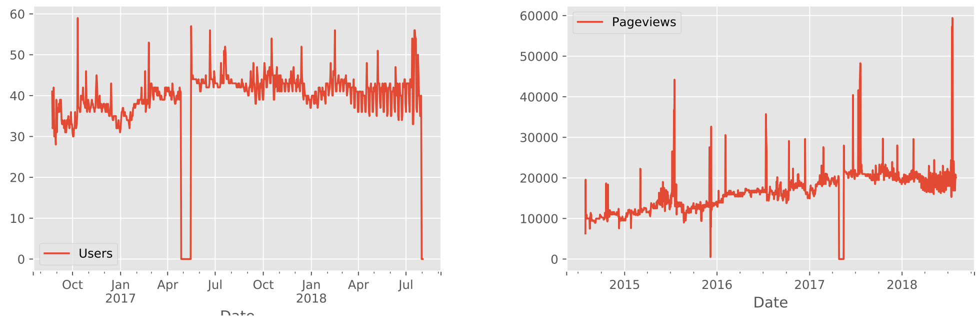

(**a**) Number of referrers since start of trial.

(**b**) Pageviews since the start of trial.

Figure 6.1: Pheme statistics from the start of the trial using the *Google Analytics Injector*.

enables us to collect and process large quantities of analytics data from the e-Campus display test-bed.

In particular, the injection module has been designed to map and inject incoming signage-specific analytics data to Google Analytics in real time, i.e. to perform the mapping to a compatible data model and inject the data stream immediately via HTTP requests. The modularity of the Pheme system and the existence of appropriate base classes enabled the development of the injection module with minimal effort. Listing 4 shows a code snippet of the Google Analytics injector consisting of 23 lines of Python code.

The injector extends `HTTPPostInjector`, a base class implementing HTTP Post requests to any external service. The `pre_inject(self)` implemented in the child class is automatically called by Pheme for each incoming dataset which is made available in `self.data`. The `pre_inject(self)` method is required to return the mapped dataset in the form of a Python dictionary which are automatically injected into the third-party system using a standard HTTP request. In this case, the `HTTPPostInjector` implemented a simple HTTP Post request to the external service with the mapped dictionary in the payload. The URIs of the target API of the external service ar specified within the `pre_push(self)` method.

6.2.3 Reported and Captured Analytics Data

The Pheme trial was initiated in August 2014 and has been in daily operation for over four years. In this time period, Pheme has collected and processed a total of 159,264,530 analytics-related events yielding a database of 73.67 GB (excluding indices). Pheme typically receives approximately 3.8 requests per second which typically are reported analytics events using the dedicated analytics reporting API.

Over the course of the data collection period, the number of reporting displays has increased from 21 (August 2014) to 65 (August 2018). Figure 6.1a provides an overview of the number of referrers over the past 365 days retrieved from Google Analytics (i.e. reported through Pheme's Google Analytics injector). Figure 6.1b visualises the number of daily pageviews since the start of the trial – consisting of a mean daily pageview requests of 16411.04 (Mdn: 16467.00, SD: 5291.78). The high standard deviation is likely caused by

```python
from .base_injectors import HTTPPostInjector
from .constants import google_analytics


class InjectorGoogleAnalytics(HTTPPostInjector):

    def pre_push(self):
        self.api_path = google_analytics.API_COLLECT_PATH
        self.api_url = google_analytics.API_URL

    def pre_inject(self):
        """
        Converting self.data to a data structure compatible with Google
        Analytics based on UMP. This method returns a dictionary that
        can be directly pushed to Google Analytics as JSON.
        """

        # Output dictionary
        ga_data = dict()

        # Mapping of general analytics metadata
        ga_data['cid'] = self.data.cid
        ga_data['tid'] = self.get_linked_tid()
        time_difference = timezone.now() - self.data.created
        ga_data['v'] = google_analytics.PROTOCOL_VERSION
        ga_data['t'] = hit_type
        ga_data['ua'] = self.data.request_user_agent

        if hit_type == HIT_TYPE_PAGEVIEW:
            # Pageview specific
            ga_data['dl'] = self.data.dl
            ga_data['dh'] = 'http://signanalyticsproject.appspot.com'
            ga_data['dp'] = self.data.dp
            ga_data['dt'] = self.data.dt or "None"
            ga_data['cd'] = self.data.cd or "None"

        # Mapping for other hit types can be inserted here.

        return ga_data

    def __str__(self):
        return "Google Analytics injector to report display content changes."
```

Listing 4: Pheme example injector implementation to support Google Analytics.

Table 6.1: Number of reported content requests to Pheme over the 365 days by content type.

Content Type	Request Count	Mean Duration [s]	SD [s]
All	2376	21.60	25.97
Image	2098	16.77	4.11
Stream	9	80.33	112.97
PDF	54	16.07	1.66
Video	155	87.30	67.87
Website	53	16.47	6.98
Other	7	19.44	3.34

the increase of displays over time. Similarly, the constantly increasing number of reported pageviews is a result of the growing display test-bed. The figures show the high stability of Pheme and the Google Analytics injector with only two noticeable drops of incoming requests throughout the entire trial (at the end of 2016 and in July 2017). Both instances were a result of multi-day power cuts caused by flooding (December 2016) and damages caused by construction work (July 2017) causing the e-Campus test-bed to shut down.

To further analyse the response and accuracy of reported content requests and injections into Google Analytics, we consider the number of content items, content times and resulting average play times for content items reported from displays via Pheme to Google Analytics. Due to the data retention policy of Google Analytics, detailed insights (e.g. logs including referrer, reported content views and average time spent for each content item) can only be retrieved for a period of one year. Therefore, to analyse the accuracy and plausibility of reported content through Pheme to Google Analytics we consider logs provided in the past 365 days (Table 6.1). In total, displays reported to Pheme 2376 unique content items which were visible on displays on a mean of 21.60 seconds (SD: 25.97). We categorised the content into six distinct content types: *image*, *stream*, *video*, *website* and *other*. Typically, e-Campus displays have been configured to show static content for a duration between 15-20 seconds whilst videos and streams are shown for the duration of the video – leading to a higher mean play time as well as higher standard deviation due to different video lengths across the content items. In contrast, content items of type image have a low standard deviation of the content playout time suggesting consistent durations – reflecting the case and configuration of e-Campus displays.

6.2.4 Example Reports from Pheme

We note that a set of example reports that can be computed based on the Pheme dataset and utilising the Google Analytics injector have been introduced in Section 4.3 (Using Web Analytics Engines for Display Analytics Reporting, p. 102).

6.3 Tacita: Client-based Tracking

In this section, we evaluate the user experience, accuracy of the system data for public display analytics and provide system benchmarks for Tacita. The system was initially introduced in Section 3.4.1 (Viewer-based Tracking, p. 66) as a mechanism to collect viewer-centric analytics data about passers-by whilst at the same time allowing viewers to personalise the content on nearby displays. Due to the large amount on user interaction and mobility traces data collected, the underlying dataset can be utilised to create analytics reports on viewer interactions with displays as initially described in Section 4.2.5 (Display Personalisation Retention Analytics, p. 98).

Excerpts of this section are based on the following publication:

- Mateusz Mikusz, Peter Shaw, Nigel Davies, Sarah Clinch, Ludwig Trotter, Ivan Elhart, Marc Langheinrich, and Adrian Friday. "Experiences of Mobile Personalisation of Pervasive Displays". In: *ACM Transactions on Computer-Human Interaction – TOCHI (in preparation)* (2018)

6.3.1 Methodology and Datasets

6.3.1.1 Integration in the Context of e-Campus

In order to allow us to conduct a long-term and in-the-wild trial of Tacita, we integrated the system into the e-Campus display test-bed and deployed it as a service to students, staff and visitors across the University campus. In order to fully support Tacita, substantial additions were required to the previously introduced systems and components including the Lottery Scheduler, particularly regarding the interactions of displays with the Display Gateway and associated interfaces (Figure 6.2, highlighted in blue). Figure 6.3 provides an overview of the communication and data flow between the Display Gateway and subsequent components within the Lottery Scheduler.

In the subsequent section, we provide a brief overview of the additions and implementations of individual system components.

Display Gateway Interfaces

The Display Gateway was initially introduced in Section 3.4.1.1 and serves as the application programming interface for third-party applications requesting screen time, e.g. based on viewer presence at a particular location in the vicinity of the display. The Display Gateway communicates with individual display nodes through a two-way communication channel enabling displays to receive and process content scheduling requests in real time. We modelled the interface for communicating with the Display Gateway as a virtual sensor within Yarely's Sensor Management component (Figure 5.4). Similar to other components in Yarely, the sensor has access to the internal communication channels and the Context Store of the Lottery Scheduler (Figure 6.3). The Display Gateway includes the description of the request time

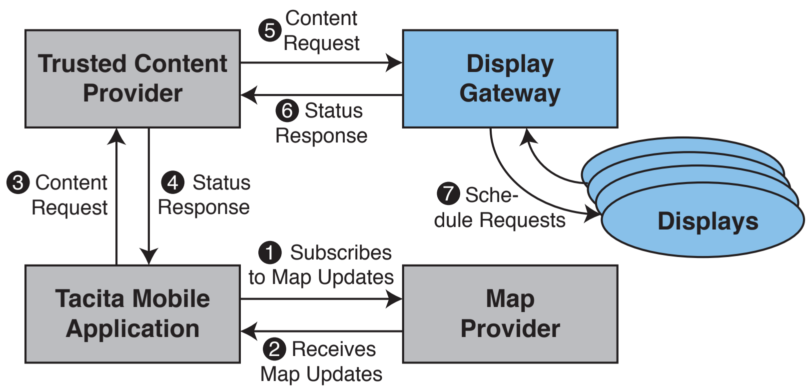

Figure 6.2: Tacita system architecture focussing on the Display Gateway and interfaces on the display node.

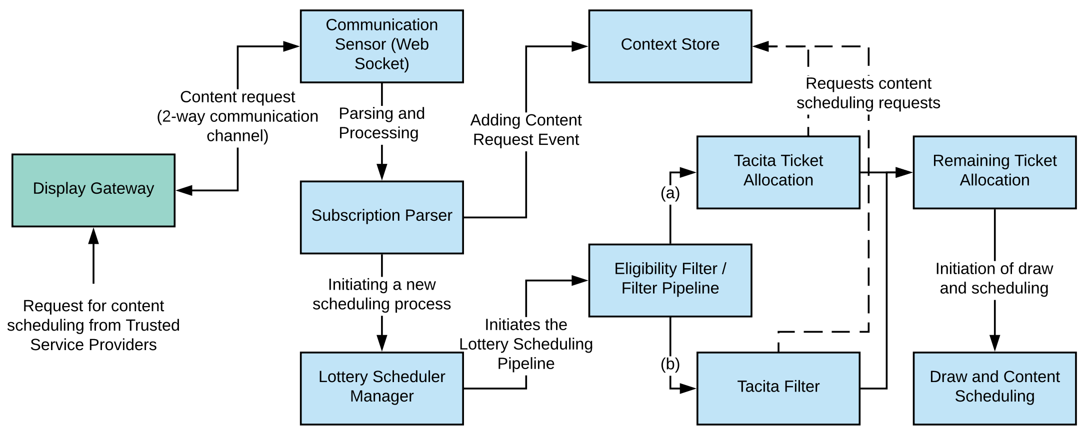

Figure 6.3: Interaction and request flow of Tacita requests across the Display Gateway and Yarely.

and the Content Descriptor Set of the requesting third-party application in the payload of the request. The sensor validates incoming messages (e.g. ensuring the trusted source and expected format) and forwards the request, similar to subscription updates, to the Context and Constraints Parser for further processing using the internal eventing system of Yarely. The Context and Constraints Parser processes the CDS by parsing and transforming it into an internal object-based format and stores it within the Context Store for further consideration. Due to the use of the CDS format, we support the full stack of features as part of the CDS such as nested content items (i.e. content that is composed of multiple items), content scheduling requirements and constraints (e.g. content date and time availabilities).

Filtering and Ticket Allocation

After parsing and storing requests in the Context Store, the Context and Constraints Parser module initiates a new Lottery Scheduling process to allow the filtering and lottery allocation components to immediately consider and react to changes in the context of the display. The lottery scheduler has the opportunity to re-evaluate the previous content decision and determine whether the currently shown item is still appropriate. In order to support Tacita, we have considered the following two approaches (Figure 6.3): supporting Tacita via (a) the *Tacita Ticket Allocator*, and (b) the *Tacita Filter Module*. In the case of approach (a), the probability for the requested item is dependent on the number of lottery tickets that have been made available for the Tacita Ticket Allocator in relation to other potential ticket allocators, and the number of tickets the ticket allocator has associated with the requested content item. In the case of approach (b), the Tacita Filter has already removed any content that has not been requested by a client – leaving the remaining set of eligible content items with requested content only. Therefore, the lottery process will yield the requested content item in any case (e.g. if multiple content items have been requested, one content will be determined at random). This approach is particularly useful to support walk-by personalisation in which case we ensure that passers-by will *always* see the requested content under the assumption that a content scheduling request has been successfully issued. To maximise the user experience and ensure that personalised content is provided to the user as often as possible, we have chosen to implement approach (b) for the entire duration of the trial.

6.3.1.2 Trial Context and Collected Datasets

To support the evaluation of Tacita including all components, we collected a large set of quantitative measures as part of an in-the-wild study at Lancaster University. The study duration consisted of 206 consecutive days (from May 2017) and 44 displays equipped with BLE beacon sensors. During the study, a total of $226,620$ events were captured including $24,565$ content personalisation requests from a cohort of 147 users. We only consider users who installed Tacita on their mobile device and issued at least one content personalisation request within the study, enabling us to remove users without any intentions of using the

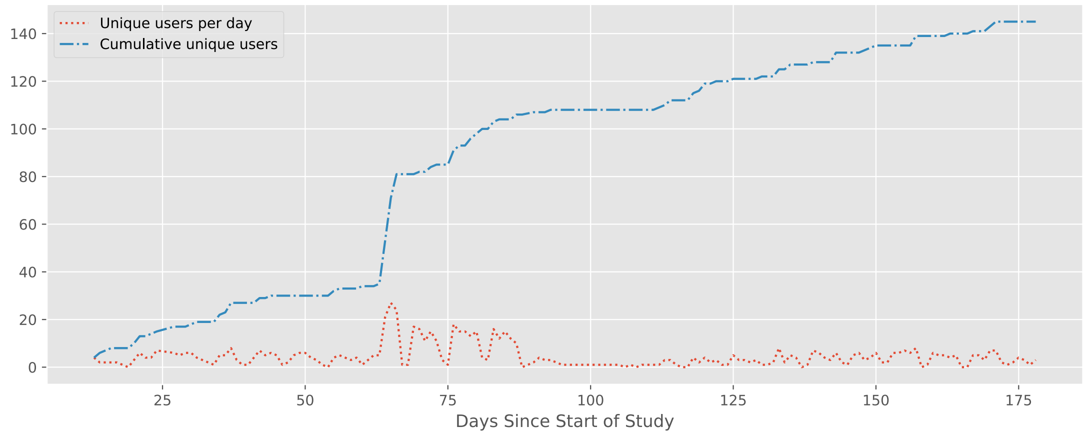

Figure 6.4: Tacita users at Lancaster University over the study period (initially published in [Mik+18d]).

system. Figure 6.4 provides an overview of the growth of the user base during the study period including fast growths during periods of active recruitment.

To enable accurate and comprehensive data collection for the subsequent analysis of Tacita, we instrumented all components of the Tacita ecosystem (i.e. Trusted Content Providers, the Display Gateways and the Tacita Mobile Client applications) to capture system and user interaction events and the following timestamps:

1. BLE beacon sightings (i.e. display proximity) on the mobile phone,
2. requests received from the Tacita Mobile Clients to Trusted Content Providers,
3. requests received from Trusted Content Providers to Display Gateways,
4. display opening and showing the content from the requested Trusted Content Provider, and
5. viewers accessing the configuration page of Trusted Content Providers through the Tacita Mobile Client.

The above mentioned timestamps have been captured as follows. Event (1) has been logged on the user's device with a timestamp at the point at which the iOS background process detected the beacon in proximity to the user's device and called the location tracking method of the Tacita Mobile Client application. In addition to the timestamps, the Tacita Mobile Client creates a unique request identifier (UUID version 4) to each beacon sighting, enabling us to trace and capture the latencies of the request through the chain of subsequent API requests throughout the Tacita ecosystem from the Trusted Content Provider to the Display Gateway and display nodes. Events (2), (3) and (4) were logged on our backend systems with the associated timestamp, and server clocks have been synchronised for those events that were executed on separate servers. We compute the latencies between system components by matching associated requests using the unique request identifier. The mapping of beacon

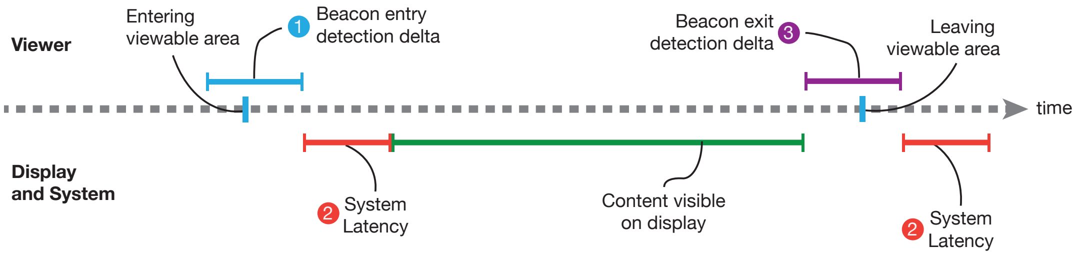

Figure 6.5: Overview of the content delivery process in pervasive display systems together with critical events affecting the proximity detection performance. Beacon entry (1) and exit detection deltas (3) depend on the underlying proximity detection technology, whereas system latency (2) depends on network and system performance (initially published in [Mik+18d]).

sighting to each display, and the subsequent analysis of beacon sightings and associated content requests have been conducted on server-side based on the requests that have been recorded on Trusted Content Providers.

In addition to system logs and timestamps, we captured user interactions with Trusted Content Providers through the Tacita Mobile Client by logging access to the configuration pages on our backend. In particular, this included capturing occurrences of users accessing the configuration page of a Trusted Content Provider and changing their preferences – including the configuration values, timestamps and anonymous user identifiers.

6.3.2 Viewer Detection Accuracy and Performance

We first conducted an analysis of the dataset to better understand (1) the accuracy of beacon sightings for display analytics purposes, e.g. for determining viewer dwell times in the proximity of displays, and (2) the latency and performance of the overall system implementation to support walk-by personalisation as one of the most challenging personalisation techniques.

6.3.2.1 Beacon Detection Accuracy and Performance

The timely detection of proximate BLE beacons is fundamental to accurately detect the viewer's proximity to displays. In particular, Tacita relies on BLE beacon detection to deliver personalised content to a viewer as they pass by a display, i.e. support walk-by personalisation. The detection of BLE beacons through the viewer's mobile devices therefore has a high impact on both the delivery of personalised content and the detection of viewer behaviour in the proximity of displays. To better understand the potential latencies and performance impact, we decompose the chain of beacon detection and API calls into individual steps and measure each step of the Tacita chain in detail. We illustrate the complexity of system components involved in viewer detection and potential errors affecting the detection accuracy and overall system performance in Figure 6.5:

(1) Beacon Entry Detection Delta specifies the time delta between the viewer entering the proximity of a display (and the adjacent BLE beacon) and the time the background location tracking detects the proximate beacon to send the request to the infrastructure.

This delta can be positive if the proximity was detected *after* the viewer has entered the viewable area of the display and negative if the proximity was detected *before* the viewer entered the viewable area of the display.

(2) System Latency specifies the subsequent latencies that occur due to the API calls to Tacita system components. In particular, this includes the API call from the viewer's mobile device to the Trusted Content Provider, and subsequent calls to the Display Gateway and the display node in proximity to the viewer.

(3) Beacon Exit Detection Delta specifies the detection of a viewer leaving the range of a beacon. Similarly to the entry detection delta, the delta for leaving a beacon range can be negative if the detection takes place *before* the viewer leaves the area and positive if the detection takes place *after* the viewer leaves the area.

The accurate determination of viewer behaviour in proximity to the display is highly influenced by the above mentioned latencies. In particular, latencies (1) and (2) impact the accurate detection of the viewer entering the viewable area of the screen – affecting both the timely delivery of personalised content to the display as the viewer walks up to or by the display, and the accuracy of display analytics using BLE beacons to determine when the viewer has entered the viewable area and which piece of content they may have seen. In addition, the accurate detection of viewers leaving the proximity of the display is important for the accurate computation of display analytics related metrics including viewer dwell times, and to enable the system to remove personalised content and free up display real estate when the viewer is no longer able to see the display. We note that latencies (1) and (2) are affected by a combination of the beacon transmission power, frequency in which the beacon payload is transmitted and the background and location processing of the viewer's device. We are not able to influence the operating system background tasks and libraries, however, the beacons transmissions were configured to at least cover the viewable area of the displays.

In the subsequent sections we highlight both the entry, exit and system latencies in more detail.

6.3.2.2 Prototype System for Beacon Detection

In order to better understand the influence of latencies induced by the iOS operating system libraries and background location tracking tasks for the detection of nearby BLE beacons (iBeacons), we designed and conducted a controlled experiment. We utilised a controllable BLE iBeacon that provided us with the ability to accurately control the Bluetooth transmission start and end times. As a mobile device, we utilised an iPhone 6 as one of the most common mobile devices and a prototype mobile application implementing the following two distinct beacon detection mechanisms:

Monitoring: Core Location Framework Apple provides with the Core Location Framework the ability to *monitor* for the proximity of BLE beacons in the background, i.e.

with the mobile device in standby. In accordance to the Apple Developer Guidelines, a maximum of 20 beacons (a tuple of beacon major, minor and unique identifiers) can be registered for background location tracking – limiting the total number of supported beacons presumably to protect the resources of the user's mobile device. The developer is only required to implement handlers for beacon entry and exit events which are called by the operating system. Fine-grained controls over the tracking frequency are not possible. We note that for beacon detection, Apple describes the Core Location Framework as the recommended implementation way – however, it requires the user to permit 'background location tracking' for the mobile application utilising this framework. The Core Location Framework also represents the technique used in the Tacita Mobile Client. We note that in the subsequent sections we refer to this technique as *monitoring*.

Ranging: Core Bluetooth Framework The Core Bluetooth Framework allows developers to *range* for proximate Bluetooth devices – including BLE beacons. This method requires the device and the application using this technique to be active and in the foreground. In contrast to monitoring, using the Core Bluetooth Framework allows the specification of ranging frequencies and therefore provides us with the ability to develop custom entry and detection algorithms. We note that the subsequent sections we refer to this technique as *ranging*.

We captured timestamps on the machine running the controllable beacon of the transmission start and end, and on the mobile device client of the subsequent beacon entry end exit detection events that were transmitted to a virtual Trusted Content Provider for logging purposes. Both the mobile device and Trusted Content Provider were located on the same local area network minimising network-related transmission delays. We performed 10 repetitions for beacon entry and exit detection each. For the monitoring mode, we compare two states of the mobile device: active (i.e. with the screen turned on and the phone unlocked) and inactive (i.e. with the screen turned off and the mobile device on standby). For the ranging mode, we only use the mobile device in active mode as active ranging in background is not permitted as per iOS developer guidelines.

Upon determining the latencies of beacon detection times on the mobile device, we designed a follow-up experiment in a realistic setting investigating the influence of physical layouts on the transmission and detection accuracy. We designed the experiment based on the typical walk-by scenarios in which the content needs to change in time for the viewer before passing by the display and therefore allowing to capture accurate analytics in a timely manner. We performed this experiment by identifying a representative display deployment within the e-Campus displays network. The display was equipped with a single BLE beacon, and typical hardware (Mac Mini late 2014, Intel Core i5 processor with 2.6 Ghz, 8GB memory and macOS 10.10.3), and the experimenter equipped with the similar prototype application and mobile device described previously. The display was located in the foyer area of an office

(a) Picture of the experimentation area annotated with the three routes.

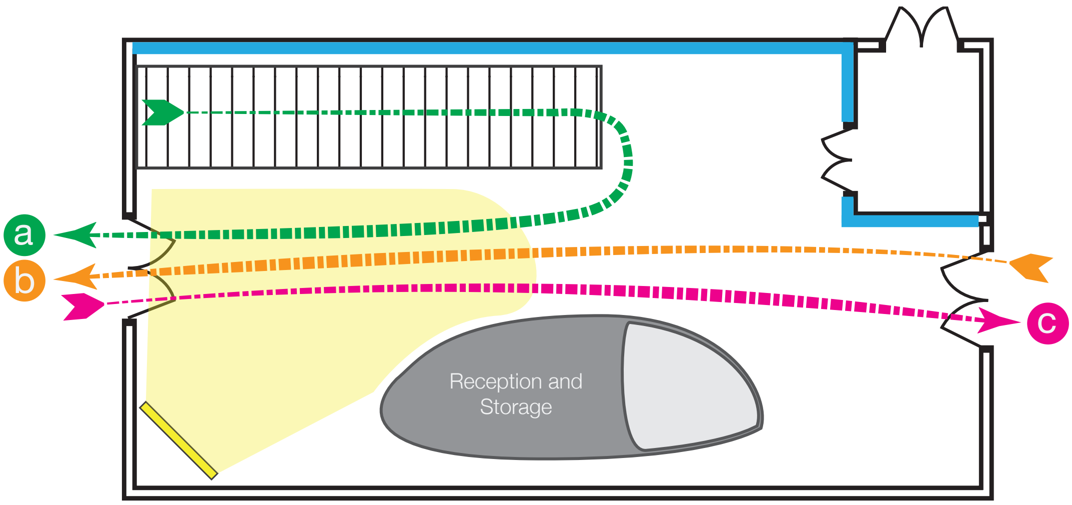

(b) Layout (not to scale).

Figure 6.6: Floor plan of the controlled walk-by experiments to capture beacon entry and exit detection latencies (initially published in [Mik+18d]).

Table 6.2: Median, mean and standard deviation for enter region (beacon detected) and exit region (beacon lost) events (seconds). We note that the ranging functionality in iOS is only available with the phone in *active* state (initially published in [Mik+18d]).

		Enter Region			Exit Region		
Condition	Phone State	Median	Mean	SD	Median	Mean	SD
Monitoring	standby	2.0	3.11	2.48	29.25	28.73	2.01
Monitoring	active	0.57	0.88	0.84	30.05	29.78	0.86
Ranging	active	0.73	0.71	0.30	10.37	10.33	0.27

building (Figure 6.6). As shown in Figure 6.6, we identified three typical ways in which viewers approach and pass by the display:

(A) The viewer approaches the display from upstairs through an open staircase introducing the difficulty of detecting the viewer as passing through floors,

(B) The viewer walks towards the display on the same floor – representing the most common form of walk-by personalisation in our display deployment, and

(C) The viewer walks towards the display on the same floor from a starting point that is separated from the display through a concrete wall.

We conducted walk-by experiments using our prototype application (i.e. capturing insights for both ranging- and monitoring-based detection techniques), capturing the following timestamps:

1. the viewer entering the viewable area of the display, i.e. the first opportunity the display can be seen (visualised in yellow in Figure 6.6),
2. the mobile device detecting the proximate beacon (entry event), i.e. the earliest time at which the system is able to react to the viewer,
3. the viewer leaving the viewable area of the display, and
4. the mobile device detecting that the viewer has left the proximity to a beacon (exit event).

For each of the three routes, we conducted 10 repetitions with the mobile device in active mode enabling both ranging and monitoring and ensuring that a WiFi connection was established at all times.

6.3.2.3 Beacon Entry and Exit Detections – Controlled Lab-based Experiment

We first consider the results of the stationary experiment in which we illustrate entry and exit detection latencies induced by the operating system and potential background processes. The resulting latencies in detecting entry and exit are shown in Table 6.2.

Table 6.3: Median, Mean, and standard deviation for beacon entry detection from entering and leaving the viewable area of the display respectively (initially published in [Mik+18d]).

		Enter Region			Exit Region		
Route	Condition	M [s]	Mdn [s]	SD [s]	M [s]	Mdn [s]	SD [s]
(A)	Ranging	3.19	2.73	3.66	10.22	10.50	1.85
(A)	Monitoring	3.66	−1.01	15.72	26.32	31.65	10.68
(B)	Ranging	5.58	5.48	3.7	10.18	10.30	1.59
(B)	Monitoring	5.10	4.89	3.24	31.37	43.89	39.38
(C)	Ranging	1.65	0.90	1.95	12.07	11.56	2.96
(C)	Monitoring	−1.23	−0.2	7.21	33.14	36.35	11.42

The entry detection (i.e. simulating the case in which the viewer enters the proximity to a display) performs well across both background location tracking and Bluetooth ranging with the device in active mode. However, the standard deviation for the background location tracking technique is slightly higher suggesting the potential impact of other background processing tasks running on the operating system yielding Bluetooth ranging as the most reliable and stable beacon detection technique. In contrast, the device in standby mode yields a noticeably higher standard deviation for detecting beacon entry with a median of approximately 2 seconds. Considering the average walking time of ≈ 1.4 m/s, we note that using Bluetooth ranging and location tracking as detection modes lets the viewer move approximately 1-5 meters before their proximity to the display is detected depending on the mode of their mobile device. Providing the typical range for BLE beacons (and Bluetooth transmission in general) is approximately 10 meters, the latency would still allow the system to react to the viewers presence in the viewable area fast enough to change the content in time – providing that the beacons have been configured with an appropriate signal strength.

The exit detection (i.e. the viewer leaving the viewable area of the display) performed significantly worse for both monitoring and ranging as a technique compared to the entry detection with a median delay of 29.25 seconds (SD: 2.01) and 30.05 seconds (SD: 0.86) with the device in standby and active modes respectively. We believe that this is a result of iOS treating background location tracking for leaving areas with either a lower performance, or applying larger thresholds before an exit event has been sent to the client application. Using ranging as a technique, we were able to achieve a significantly better and more stable performance with a median of 10.37 seconds (SD: 0.27). Due to the relatively stable exit detection (low standard deviations), we believe that the event is nevertheless suitable for display analytics purposes. However, if the display infrastructure relies on prompt and accuracy exit events to free up display real estate and remove personalised content, both monitoring and ranging may be unsuitable.

6.3.2.4 Beacon Entry and Exit Detections – Controlled Walk-by Experiments

After measuring potential delays caused by the operating system and system libraries, we now highlight the results from our controlled walk-by experiments that enable us to measure potential delays caused by the spatial layout and other intrusions. The resulting entry and exit latencies for the three routes are shown in Table 6.3.

Taking beacon detection into more realistic settings we observe highly variable results depending on the walking path and detection technique used. Similarly to the results observed in the controlled lab-based experiments, ranging yielded more consistent latencies with a significantly lower standard deviation for both entry and exit detection across all routes whilst monitoring performed faster considering entry detection only with means of -1.01, 4.89 and -0.2 seconds for routes A, B and C respectively. Both routes A and B originate in closer proximity to the display which is likely the reason for the better performance compared to route B which originates further away and simulates the case of a viewer walking towards the display. The higher standard deviation for monitoring is likely the cause of unknown background processes of the operating system and radio modules which may have an impact on the detection times. Additionally, we observe negative means in the monitoring technique, i.e. the detection of BLE beacons prior to the user entering the viewable area of the display.

In contrast to the entry detection latencies, we observe a significant improvement of using ranging instead of monitoring for exit detections. In addition to a lower delta for detecting viewers exiting the viewable area of a display, we are also able to observe a lower standard deviation yielding a better consistency in the detection times. The mean detection time using monitoring as a technique was captured at 31.65 (SD: 10.63), 43.89 (SD: 39.38) and 36.35 seconds (SD: 11.42) for routes A, B and C respectively. In contrast, we were able to capture mean delays using ranging at 10.50 (SD: 1.85), 10.30 (SD: 1.59) and 11.56 seconds (SD: 2.96). Fast and reliable exit detection is important for both freeing up display real estate once the viewer has left the viewable area and to capture display analytics with a high confidence. We believe that whilst the current state-of-the-art technology for detecting viewer proximity to BLE beacons (i.e. the monitoring technique) is not yet suitable for accurate display analytics purposes, investigating ranging as an alternative highlights the opportunities for potential future improvements to the technology.

We note that entry and exit detection latencies are highly dependent on (1) the spatial layout of the area in which the display and BLE beacon have been placed, (2) potential background task and radio processing on the viewer's mobile device, and (3) the technology used to detect viewer proximity. For example, the mean detection for route B using monitoring as a detection technique (4.89 seconds) allows the viewer to walk ≈ 6.5 meters before successfully being detected in the proximity of the display. BLE beacons therefore need to be configured accordingly with appropriate transmission power and frequency to ensure a fast enough detection for delivering personalised content to viewers walking by. Using such display sightings to calculate viewer-centric display analytics highlights the currently low accuracy and reliability of the dataset, particularly for detecting viewers who leave the viewable area of displays – a crucial insight for calculating dwell times and other analytics metrics.

Table 6.4: Median, mean and standard deviation (seconds) of the latencies between Tacita system components (initially published in [Mik+18d]).

System Components	Median [s]	Mean [s]	SD [s]
Mobile to Content Provider	0.09	2.15	16.52
Content Provider to Display Gateway	0.02	0.19	5.82
Display Gateway to Content Show	1.33	3.16	7.19
Mobile to Content Show	1.66	4.8	11.75

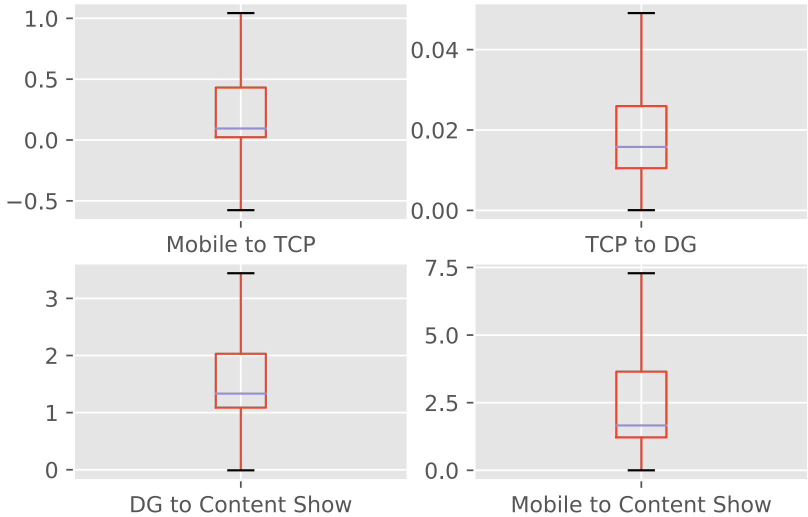

Figure 6.7: System response latencies [seconds] for the chain of Tacita system components (Mobile: Tacita Mobile Client; TCP: Trusted Content Provider; DG: Display Gateway; Content Show: requested content shown on the nearby display) (initially published in [Mik+18d]).

6.3.3 System Component Latencies

In addition to the latencies resulting from detecting proximate BLE beacons, we further investigated system-related latencies that may occur when proximity sightings and content requests are passed through the Tacita system components. As shown in Figure 6.5, such system- and network-related latencies are particularly crucial to ensuring that personalised content is delivered fast enough to an individual passing by a display. The assignment of a random globally unique identifier to each request issued by the viewer's mobile device (i.e. at the point at which a proximate BLE beacon was detected), we are able to trace and time each request throughout the chain of APIs and system components.

An overview and aggregation of the request timings captured through our in-the-wild study are shown in Table 6.4 and Figure 6.7. In particular, we consider the following latencies:

Mobile Client to Trusted Content Provider describes the latency from the point at which the mobile phone detects a proximate BLE beacon to the point at which the request has arrived at the Trusted Content Provider. The Tacita Mobile Client transmits a timestamp of the time at which the beacon was detected, allowing us to capture potential latencies that include transmission delays relating to the mobile network.

Trusted Content Provider to Display Gateway describes the latency from the point at which the Trusted Content Provider issues content requests that arrive at the Display Gateway.

Display Gateway to Content Shown describes the latency from the point at which the Display Gateway issues a request to the display node, and the display node shows the content on the display (captured through logs on the Trusted Content Provider). In particular, this latency includes potential processing and display scheduling latencies that may occur on the display itself. If a display has not shown the requested content, we consider this as a *failed request* in the subsequent analysis.

Overall Latency – Mobile Client to Content Shown describes the overall latency from the point at which the Tacita Mobile Client detects a proximate beacon and requests the content to be shown to the time at which the display shows the requested content.

Considering the overall latency, we were able to capture a median of 1.66 and mean of 4.8 seconds (SD: 11.75). These measures highlight that just considering system and networking performance, Tacita is fast enough to respond to content scheduling requests and support walk-by personalisation. In addition to the network performance, the previously discussed latencies for detecting beacons on the mobile device need to be further considered with regards to configuring the beacon range accordingly.

Considering the latencies between individual system components, we firstly observe a median latency of 0.09 seconds (mean: 2.15, SD: 16.52) between detecting a proximate BLE beacon and the request arriving at the Trusted Content Provider. Whilst for the majority of cases this latency is low, the high variance suggests the potential impact of poor mobile data connectivity and network connection. In particular, such issues may arise when viewers transition between in- and outdoor locations or the mobile device switches from a cellular data connection to Wi-Fi. We measured the delays between Trusted Content Providers and Display Gateways at a median of 0.02 seconds (mean: 0.19, SD: 5.82) as the lowest latencies in the chain of system components. This is likely due to the negligible network latency (both systems were located on the same virtual network) and short processing and parsing times. The second highest latency was identified between the Display Gateway sending the request to a display node and the display node showing the requested content (median: 1.33, mean: 3.16, SD: 7.19 seconds). Despite the Display Gateway keeping an open communication channel with each display node, the high latency is a result of completing the scheduling process on the display node itself. In particular, for each incoming personalisation requests, display nodes have to issue a new lottery scheduling request to enable a dynamic change of content creating a noticeable processing overhead. We note that there is a potential to improve this latency by simplifying the scheduling process on the display node and, for example, always prioritising personalisation requests.

In order to measure the reliability of the overall Tacita system, we consider the total number of daily content requests and those requests that have been issued by a Tacita Mobile

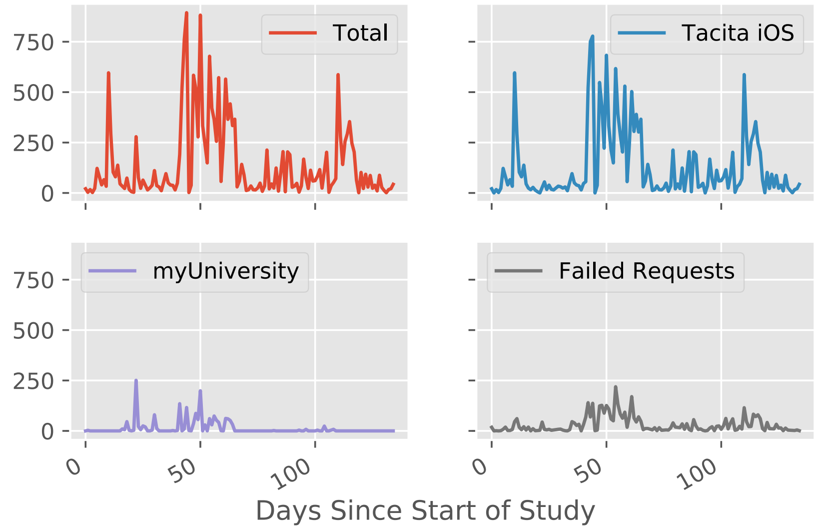

Figure 6.8: The total number of daily requests issued to Trusted Content Providers across all client applications (top left); originating from the iOS-based Tacita Mobile Client (top right); originating from iLancaster (bottom left); and the number of daily failed requests (bottom right) (initially published in [Mik+18d]).

Client but have not yielded the personalised content to be shown on nearby display nodes. Figure 6.8 provides an overview of the total number of request counts, requests issued by the individual mobile client applications (iLancaser and the Tacita Mobile Client), and the number of failed requests (i.e. requested content has not appeared on the display). We observed that Tacita was able to handle large amounts of daily requests (peaked at over 750 on certain days) of which the majority originated from the Tacita. The number of failed requests over the entire study was measured at approximately 20% of the entire amount of issued requests. However, we note that requests can fail for various reasons. For example, if a viewer has requested multiple Trusted Content Providers to show on a proximate display, at most only one of these will be scheduled to be shown on the display – leaving the remainder of requests to fail. Additionally, proximate display nodes may consider different scheduling requirements and contextual events that overrule the incoming personalisation request from the Display Gateway.

6.3.3.1 Accuracy for Analytics Data Capture

Based on the statistics reported previously, we gained a high level of insights into the performance and accuracy of using viewer-based tracking technology in the form of BLE beacons to detect viewers in the proximity of displays and measure their dwell times. In particular, we observed that viewers are detected entering the viewable area of displays with only a small delay (using monitoring as a technique between -1.23-5.10 seconds; Table 6.3). Whilst this is often fast enough to support walk-by personalisation, it also provides an accurate measure for display analytics purposes to, for example, capture the content that viewers may have seen when walking by displays, or to simply count the number of (unique) viewers across a network of displays.

Table 6.5: Details of application adoption showing the percentage of total users who issued at least one content request to the Trusted Content Provider and the availability of the Trusted Content Provider during the study period (initially published in [Mik+18d]).

Category	Trusted Content Provider	Availability (%)	Users (%)
Transport and Navigation	Bus Departures	100	65
News and Information	Weather	100	40
News and Information	World Clock	100	34
News and Information	News	100	23
Entertainment	Live TV	100	21
Social Networks	Twitter News Feed	100	8
Entertainment	Pictures	86	7

In contrast, using viewer-based tracking via BLE beacons to compute dwell times, i.e. a metric that is highly reliant on an accurate detection of viewers leaving the viewable area of a display, has proven to be challenging. Using monitoring as a technique (representing the current state-of-the-art for tracking viewers whilst their phone is in standby mode), the captured delays in our controlled experiment ranged from 26.32-33.14 seconds (Table 6.3) depending on the route. In particular, the standard deviations measured ranging across 10.68-39.38 seconds (Table 6.3) show the high variability in the results making it almost impossible to provide an accurate measure.

6.3.4 Usage Pattern Analysis

In addition to the purely systems- and performance-focussed evaluation of Tacita, we additionally investigate the usage patterns across the entire deployment.

6.3.4.1 Trusted Content Provider Usage

During the lifetime of the trial, we had a total of seven Trusted Content Providers across four categories available to any user (Table 6.5). In order to understand better the types of applications users favoured, we focus entirely on daily content requests as a metric – allowing us to filter out users who only explored the Tacita Mobile Client and configuration pages without using and requesting Trusted Content Providers to be shown on displays. We note that whilst the majority of Trusted Content Providers were available throughout the entire trial, the Pictures application was only available for 86% of the lifetime of the trial. Considering the number of unique users per Trusted Content Provider, we are able to identify Bus Departures, Weather and World Clock as the most commonly used Trusted content Providers whilst Live TV, Twitter News Feed and Pictures were the least requested content. In particular, the Twitter News Feed is the only example of a social network related application indicating that these category of content are the least interesting to the audience. Considering the total number of daily requests per application (shown in Figure 6.9), we can observe similar patterns in which the most frequently requested content include Bus Departures, Weather and World Clock. We

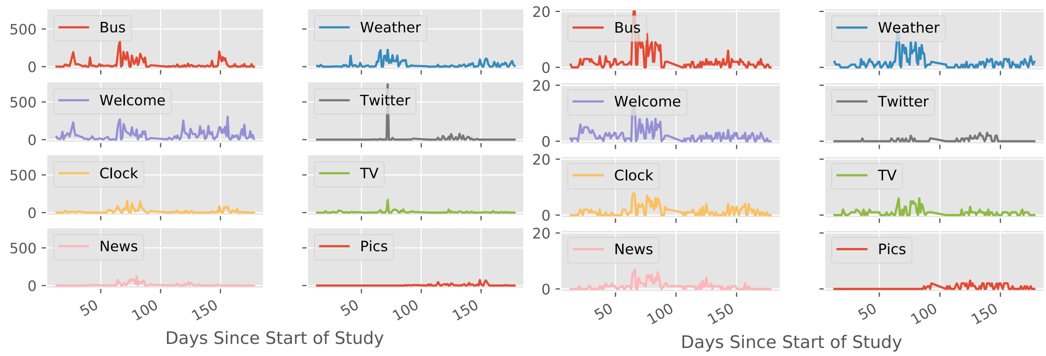

(a) Total number of *requests* per day.

(b) Total number of *unique users* per day.

Figure 6.9: Number of daily requests for each available Trusted Content Provider (initially published in [Mik+18d]).

note that due to the preselection of Welcome for any new Tacita user, we exclude this Trusted Content Provider from further analysis and discussion.

Considering the number of daily requests (Figure 6.9a) and daily unique users per Trusted Content Provider (Figure 6.9b) we can observe similar usage patterns regarding the most and least commonly requested content. In addition, the figures reflect the periods in which the Tacita service was actively advertised and users opportunistically recruited across the University campus (62 days since the start of the study), and holiday periods in which the vast majority of the campus population was not present (84-114 days since start of the study).

Considering the available Tacita Mobile Clients (i.e. the native iOS-based client application and the integration into iLancaster), we observe that the majority of users utilised the native iOS clients – as shown in Figure 6.8. Further, the majority of beacon sightings were issued automatically with the Tacita Mobile Client in the background while only a small number of requests (324 in total, 1.4%) were a result of a manually triggered content request. In total, 65 users requested content manually whilst being present in the proximity to a Tacita-enabled display of which the majority of the requests were captured on the first day of usage. Only a subset of 10 users continued to manually request content on subsequent days highlighting the importance for supporting walk-by personalisation that requires only minimal interactions by the user to foster long-term and persistent usage.

A clear challenge in the deployment of display personalisation systems is the determination of conflicting requests on individual displays, i.e. how to handle multiple viewers requesting different content on the same display and therefore competing over screen real estate. In the context of the Tacita trial, however, we observed only a very small number of conflicting requests, i.e. requests from at least two individual users that arrived to any Trusted Content Provider for the identical display within a 30 second time window (the average time personalised content remains visible on a display). We observed a mean of 0.4 competing requests per day (Mdn: 0.0, SD: 0.07) illustrating the currently low significance of this issue.

Figure 6.10: Numbers of requests of Trusted Content Providers per display location (green: low number of requests; red: high number of requests) across all applications (initially published in [Mik+18d]).

Figure 6.11: Dwell times of viewers in front of displays (green: low dwell times; red: high dwell times). The dwell times have been normalised based on the displays with the highest and lowest dwell times respectively (initially published in [Mik+18d]).

6.3.4.2 Spatial Request Patterns

Due to the ability of associating viewer beacon sightings and content requests to spatial locations, we are able to analyse in which locations viewers have typically requested content. This provides us with insights into the availability of the system across the entire display network, and the general viewer behaviour.

We are able to observe beacon sightings and requests to Trusted Content Providers from all of the enabled displays across the University campus. Figure 6.10 provides a heatmap of the proportions of requests across each individual display (green indicates proportionally low numbers of requests whilst red indicates proportionally high numbers of requests). We are able to confirm expected patterns of content requests: locations on the university campus that are characterised by a high volume of students and staff also yield a high volume of content requests including the library, learning areas and displays located along the main path. In contrast, we observe low numbers of requests in locations along the outskirts of the campus, departmental buildings and the university conference centre. In addition to content requests, the ability to use beacon entry and exit events allows us to compute viewer dwell times for each display location. Whilst we outlined the limitations of this metric in the previous section (particularly regarding the inaccuracy and high latency of the exit event), we used the start and exit events to compute estimated user dwell times. Figure 6.11 shows a heatmap of display locations and proportional dwell times (green indicating short dwell times, red indicating long dwell times). Displays with the proportionally highest dwell times include the university conference centre, library, learning areas and lecture theatres – reflecting the expected patterns in the context of the campus.

The insights into viewer behaviour including mobility patterns highlight that beacon sightings are a suitable technique to capture and express viewer-centric analytics reports. However, as stated previously, we note the limitations of using beacon sightings for the computation of dwell times and emphasise that the results are only indicative. Future technological advancements, however, have the opportunity to provide a higher quality and accuracy of dwell times which are a crucial metric in public display analytics. With regards to the delivery of personalised content, dwell time measures could be used to optimise content delivery for individuals. For example, if display locations and/or certain times have been identified to typically feature high dwell times, the prompt delivery of personalised content shortly after the request was issued by viewers becomes less important as the viewers are likely to remain in the viewable area of the display. In contrast, for display locations characterised by very short dwell times suggesting that viewers are walking by the display instead of dwelling inside its viewable area, the timely delivery of personalised content becomes important.

6.3.4.3 User Retention and Usage Duration

We note that we have introduced a set of analytics reports specific for display personalisation networks (including retention reports) in Section 4.2.5 (Display Personalisation Retention Analytics, p. 98). We used these reports to highlight the opportunities of novel display personalisation retention analytics. In this section, we focus on using similar types of reports to provide further insights into the validity of the Tacita trial and its results. In particular, we describe the measured retention rates and typical durations of usages of Tacita. In order to understand how long viewers have used the system, we use content requests as an indicator and measure the time periods (in days and weeks) over which we see continuous requests from single users. We note, however, that such content requests are only issued if users walk by displays, the viewer's mobile phone detects the proximate BLE beacon and a request was successfully transmitted to the Trusted Content Provider. Therefore, if content requests have not been observed for a certain duration, these may not necessarily be a cause of the user deciding to uninstall the Tacita Mobile Client or deliberately stop using the service – instead, it may be a result of the user's mobility patterns.

In order to capture a first insight into the retention and usage behaviour, we first count the number of consecutive weeks in which we observe content requests from individual users, i.e. counting the number of weeks between the first and last content request. This approach builds on the assumption that our typical user base consists of students and university staff members from which we would expect at least a single content request (i.e. BLE beacon sighting) per week. In order to ensure a valid representation of usage patterns, we only consider viewers

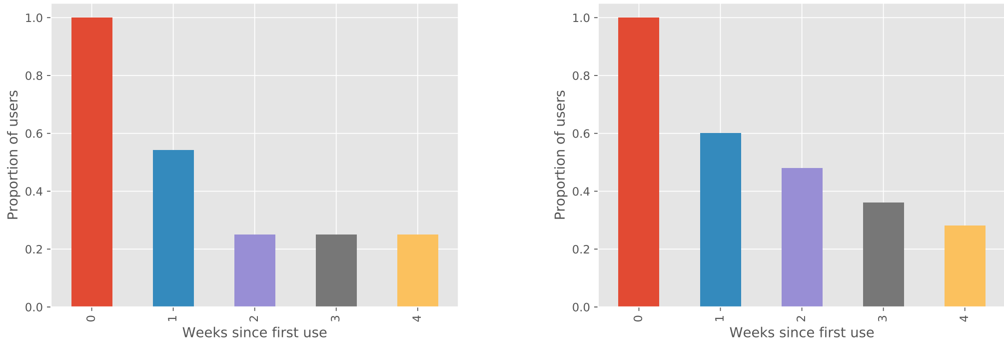

(a) Winter term

(b) Spring term

Figure 6.12: Retention weeks by term (initially published in [Mik+18d]).

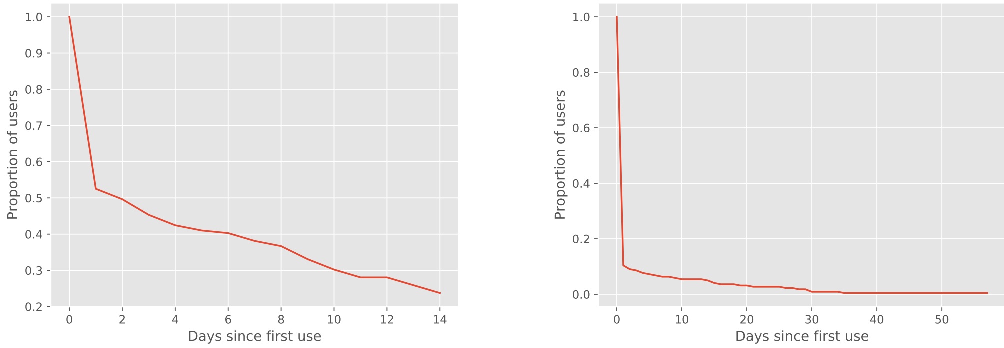

(a) Requests issued to Trusted Content Providers by Tacita users walking by displays.

(b) Revisits of configuration pages of Trusted Content Providers by Tacita users.

Figure 6.13: Retention days (initially published in [Mik+18d]).

who first started using Tacita in the first five week of the term and discard any new user who started using the service in the last five weeks of term as this set of users would have not been able to use Tacita for a total of five weeks. Figures 6.12a and 6.12b show the retention statistics for the University winter and spring term respectively. Across both terms, we observe a comparable decline in the proportions of users across the weeks. However, we also observe a small but stable set of users who continue to use Tacita over the duration of five weeks until the end of the measurement period. This finding suggests that the implementation of the service was stable enough to retain a set of users, and provide opportunities for display analytics to trace the behaviour and movement patterns of individuals over longer period of times. This finding is particularly important for potential future research such as A/B-testing of different display content and measuring the potential impact on viewer behaviour and mobility patterns across an entire space.

Figure 6.13a provides more detailed insights into the usage of Tacita over a period of 14 days. In line with the week-based retention statistics introduced above, we additionally computed mean Tacita usage durations as the delta of the first and last day content requests were issued by users. We observed that $\sim 37\%$ of users are still issuing requests after seven days, and $\sim 22\%$ of users after 14 days form the first day requests have been observed – showing that the highest fall of users takes place within the first week of use.However, as suggested by the week-based usage insights, a stable proportion of users remains to issue content requests over a period of multiple weeks. We note that there can be various reasons for the decline. Firstly, users are required to permit the Tacita Mobile Client to access their location while the mobile client is not actively used in the foreground. In iOS, an additional notification appears after using Tacita for approximately two days reminding the user about the active background location tracking feature in Tacita and giving another opportunity to deactivate the location tracking feature. Secondly, in addition to granting permission for location tracking, users must also turn on Bluetooth and activate at least one Trusted Content Provider for display sightings to be successfully reported. Further, various external factors may impact beacon detection including the viewers mobility patterns (they must pass by a Tacita-enabled display) and maintain an active network connection via Wi-Fi or cellular data. Despite this high burden of requirements, we observe high conversion and usage rates over the first weeks of the trial and are able to maintain a stable user base – allowing us to capture valuable insights about the quality and performance of the overall system.

Whilst display sightings are issued automatically upon beacon sightings and do not require explicit interactions of the user, users are required to explicitly visit configuration pages of Trusted Content Providers if they wish to apply changes. In addition to retention statistics based on observed content requests, we computed retention statistics for users explicitly accessing configuration pages of Trusted Content Providers – considering the delta between the first and last time users have accessed the configuration page. Figure 6.13 provides an overview of the proportion of revisiting a configuration page throughout the study period. We observe that 75% of users visit the configuration page once and 50% visit the configuration page a second time.

6.4 Infrastructure-based Tracking

In this section, we evaluate the suitability and system performance of Wi-Fi fingerprinting as an example of an infrastructure-based approach for tracking viewers within a defined space without the requirement for a dedicated mobile client application. This system was initially introduced in Section 3.4.2 (Infrastructure-based Tracking, p. 73). The focus of our evaluation lies particularly on gaining insights regarding the accuracy, reliability and performance of Wi-Fi fingerprinting as technique used to detect viewer proximity to public displays. We note that in contrast to Tacita, the computation and proximity detection of viewers to displays is achieved solely on the infrastructure side. This work also provides insights into the deployment of a display analytics system in the context of a large, commercial environment.

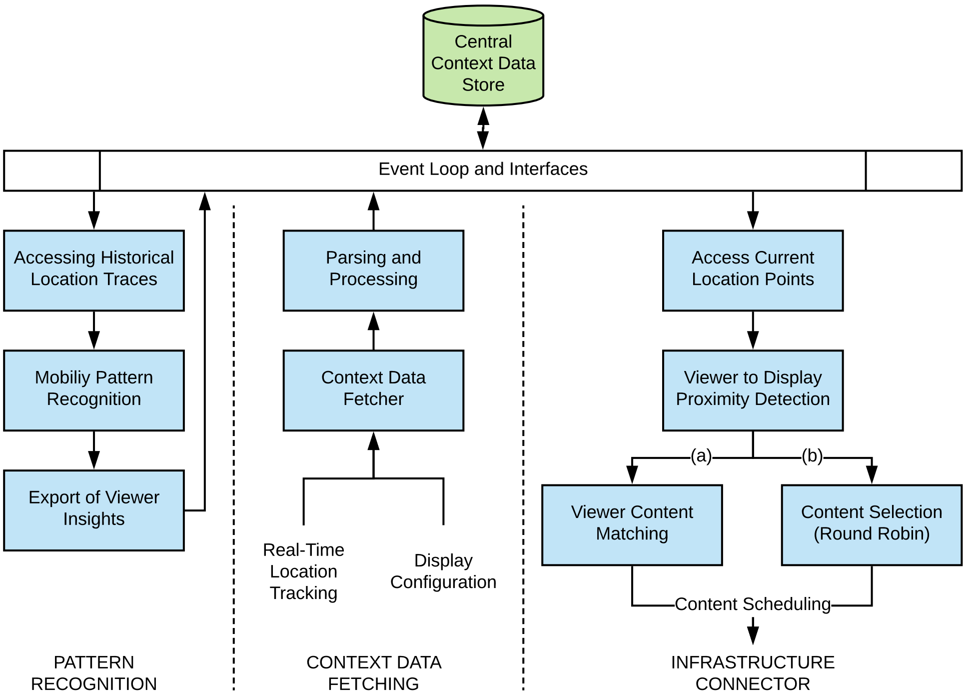

Figure 6.14: Flow diagram for infrastructure-based personalisation.

6.4.1 Methodology and Datasets

6.4.1.1 Integration in the Context of LiveLabs

The system operates in an environment that has been equipped with infrastructure-based mobility tracking technology of visitors entering and navigating within the space. Viewer mobility data is captured using Wi-Fi fingerprinting technology developed and deployed as part of LiveLabs [Jay+16]. Whilst a client device (e.g. a smartphone or tablet) is still required, users are not required to install a dedicated client application such as Tacita. By connecting to a Wi-Fi network available in the space, the location tracking provided by LiveLabs is initialised. For each device connected to the network, backend components provided by LiveLabs compute the location of the client device based on Wi-Fi fingerprinting printing. This is a contrast to the approach taken by Tacita in which the location (and display proximity) is computed by the client device itself by detecting Bluetooth Low Energy beacons. In order to conduct the trials described in the subsequent sections, we implemented the following middleware components:

1. location *pattern recognition* of viewers,
2. capturing a log of real-time viewer location information (*context data fetching*), and
3. detecting viewer proximity to displays and dynamically delivering the best-matching content to the viewer (*infrastructure connector*).

An overview of the data flow including the main components is provided in Figure 6.14. Whilst the infrastructure-based personalisation system consists of a total of six components initially introduced in Section 3.4.2 (p. 73), we specifically focus on the components that implement required interfaces to support personalisation.

Context Data Fetching

The Context Data Fetching component accesses external information and makes these available to the remainder of the system. In particular, the following set of contextual data is accessed and processed:

Real-Time Location Tracking Visitors to the experimental space are tracked using Wi-Fi fingerprinting. In order to be tracked by the system, viewers were required to connect to a Wi-Fi network accessible by LiveLabs. Whilst the location information is collected and computed on the infrastructure-/server-side, viewers' mobile devices serve as beacons as their location is determined based on the Wi-Fi connection and signal radio signal strengths. Accessing the current location information from the infrastructure-servers in real time is a crucial part of the overall system. The real-time location tracking module accesses the location tracking interface of the infrastructure in a fixed time interval of 5 seconds and fetches the current location data of all individuals. This data is provided in the form of a single comma separated values file (each line represents the current location of a user and consists of their anonymised MAC address and location point). The raw data is processed and parsed, and subsequently stored in the contextual store. The location data further serves as a foundation for subsequent processing such as pattern recognition.

Display Configuration Similar to other public display systems, we had to provide a mechanism in which display locations, trigger zones and content schedules could be defined and considered by the system. We therefore designed a simple user interface in which space owners and administrators can define required parameters including trigger zones and the default set of content items that each display is showing. The Context Data Fetcher component retrieves this set of configuration parameters in a regular interval and makes it available to the remainder of the system including the content scheduling and user to display proximity detection components.

After successfully fetching, parsing and storing new location and other contextual data, the processing module additionally distributes a 'location update' event to the main event loop, triggering the content selection and proximity detection modules. We note that additional contextual data can be processed and stored by simply creating additional modules within this component.

Pattern Recognition

The Pattern Recognition component accesses historical mobility traces stored within the internal contextual data store in a fixed time interval and performs a set of mobility pattern recognition algorithms. The mobility traces serve as a key source for gaining an understanding of viewer preferences but are not linked to individuals using their clear names or other identifiable information. Instead, mobility traces are linked to hashed MAC addresses providing a basic level of anonymity whilst still enabling the recognition of individuals returning to the space. Locations of people present in the space are captured in approximately 5-20 second intervals (depending on the processing load on the external Wi-Fi tracking system) and stored in the internal Context Data Store.

The system was designed modularly and supports the execution of multiple Mobility Pattern Recognition algorithms on the same dataset. To demonstrate the levels of insight that can be gained from analysing the location traces, we designed a simple *location scoring* algorithm. The algorithm requests mobility traces of the previous day, and maps all location points to rooms. In a subsequent step, it computes the *total time spent* per user and room. As a result, the algorithm provides a scored list of the top five rooms in which users have spent the most time throughout an entire day—and discards the remainder of the places a user visited. This result is written back into the Context Data Store and made available to other components of the system.

Infrastructure Connector

In order to support the content selection and scheduling, we developed a novel *spatio-temporal content scheduling* approach. The decision which content to show at a certain location and point of time is highly dependent on its current context, particularly the viewers present in the proximity of the display and the analytical insights gained both about the viewers. Therefore, we define the decision on which content to show as a function dependent on the viewers identity (used to retrieve their preferences), their location and the current date and time. Equally, space owners and content providers need to be able to create display schedules that define under which circumstances content is delivered to displays and viewers.

Through the definition of 'rules', we allow the specification of the content that is to be shown on displays based on three parameters: the location in which someone has spent the most time (due to the use of the simple scoring algorithm), the date and time frame for which the score was computed, and the content that is to be delivered out. We support the definition of an arbitrary amount of rules, allowing the creation of complex content schedules and supporting the consideration for viewer interests.

During a content scheduling process, the content selection components first accesses the current set of location points of anyone present in the space (Figure 6.14). In a subsequent step, the module iterates over every person's location and compares it with the predefined trigger zones of displays to detect whether an individual has been in proximity to a display. If an individual was detected in proximity of a display, the system accesses the contextual data store to find matching rules based on the above defined function (i.e. the viewers anonymised identity, the location and the current date and time). If a single matching rule was found, the back-end selects the associated piece of content to be delivered out to the display that belongs

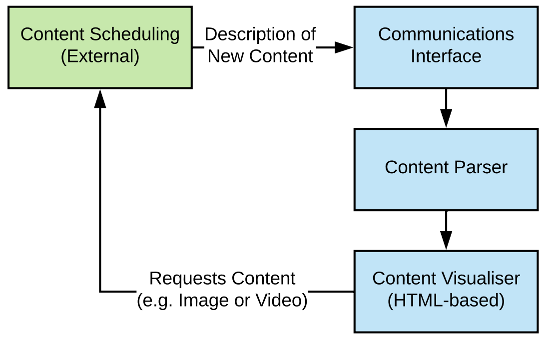

Figure 6.15: Flow diagram for infrastructure-based personalisation (display interfaces).

to the matching trigger zone. If multiple rules have been associated to an individual's context, or multiple viewers have been detected in the proximity of a single display, appropriate conflict resolution mechanisms need to be performed to determine the final piece of content to be shown on the display. A simple solution, for example, can be the selection of a rule and content at random but further, more complex algorithms for such conflict resolution can be developed.

Due to the content selection on the infrastructure-side, we a simple Web-based client is required on the display side in order to visualise the selected content items (Figure 6.15). The display client consists of an communication interface to retrieve new content scheduling commands from the infrastructure. After receiving a content scheduling request from the back-end server, the request gets passed on to the Content Parser module for parsing and verification. Content scheduling commands consist of a simple description and hyperlink to the content (in the form of a simplified version of the Content Descriptor Set) which gets passed on to the Content Visualiser module. This module retrieves the content from the remote location and shows it on the display until the arrival of the subsequent content scheduling command.

Implementation

The display client (content visualisation component for digital signs) was implemented on top of AngularJS, a JavaScript framework, and HTML. The back-end components and infrastructure connectors were implemented on top of the Tornado Web Framework and Python 3. Each component of the back-end system are implemented as individual Python-based modules, whilst content data fetchers (real-time location tracking and display configuration) are represented by Python classes, enabling a simple extension with additional context data fetching modules. Infrastructure connectors are, likewise, implemented as individual modules providing the ability to support a set of heterogeneous display infrastructures using a single back-end instance.

6.4.1.2 Trial Context and Collected Datasets

To support the evaluation of our infrastructure-based tracking mechanism, we conducted a controlled experiment in a real-world setting in the context of a large convention centre in Singapore. The convention centre spans across approximately 42,000sqm, over six stories and can hold up to 10,000 visitors. As previously described in Section 3.4.2, we utilise the LiveLabs [Jay+16] infrastructure deployed at the convention centre to retrieve current location readings. Due to the high visibility of the digital signage deployment, and limitations in the deployment of prototypical research system in the context of a commercial space, we conducted a set of controlled walk-by experiments only. Walk-by experiments allowed us to gain insights into the performance of the overall system, and the suitability of Wi-Fi-based location tracking for display analytics purposes and for the delivery of personalised content.

For designing and conducting the walk-by experiments, we followed a methodology initially described by Davies et al. [Dav+14] in which they introduced two key metrics: *content accuracy* and *content exposure* [Dav+14] defined as follows:

Content Exposure defines the proportion of time a personalised piece of content is visible on the screen while the viewer is in the viewable area of the display, i.e. the exposure of the content to the viewer. [Dav+14] describe this metric as "the effectiveness of the system at showing content to the viewer." From a display analytics perspective, content exposure is a direct result of the timely and accurate detection of a viewer in the visible area of a display and therefore fundamental to the computation of accurate analytics insights.

Content Accuracy defines the proportion of time the requested (personalised) content could have been seen by the viewer whilst the viewer was present in the viewable area of the display. Whilst this metric appears to be less relevant from a display analytics perspective as it is purely focused on content delivery aspects, capturing the content accuracy provides insights into the performance of the overall system and into the accuracy of the location tracking technique used.

We designed an experiment around the content exposure and content accuracy metrics by conducting a series of controlled walk-bys. We identified five representative display locations and walking routes within the convention centre (Figure 6.16):

Central A typical display located on the main floor of the conference venue with the viewer walking towards the display whilst remaining on the same floor level throughout the experiment.

End of corridor (1) The display is located within the main floor of the conference venue on the edge of the area that is covered by the Wi-Fi location tracking system. The viewer is walking towards the display and turning around the corner at the end of the path.

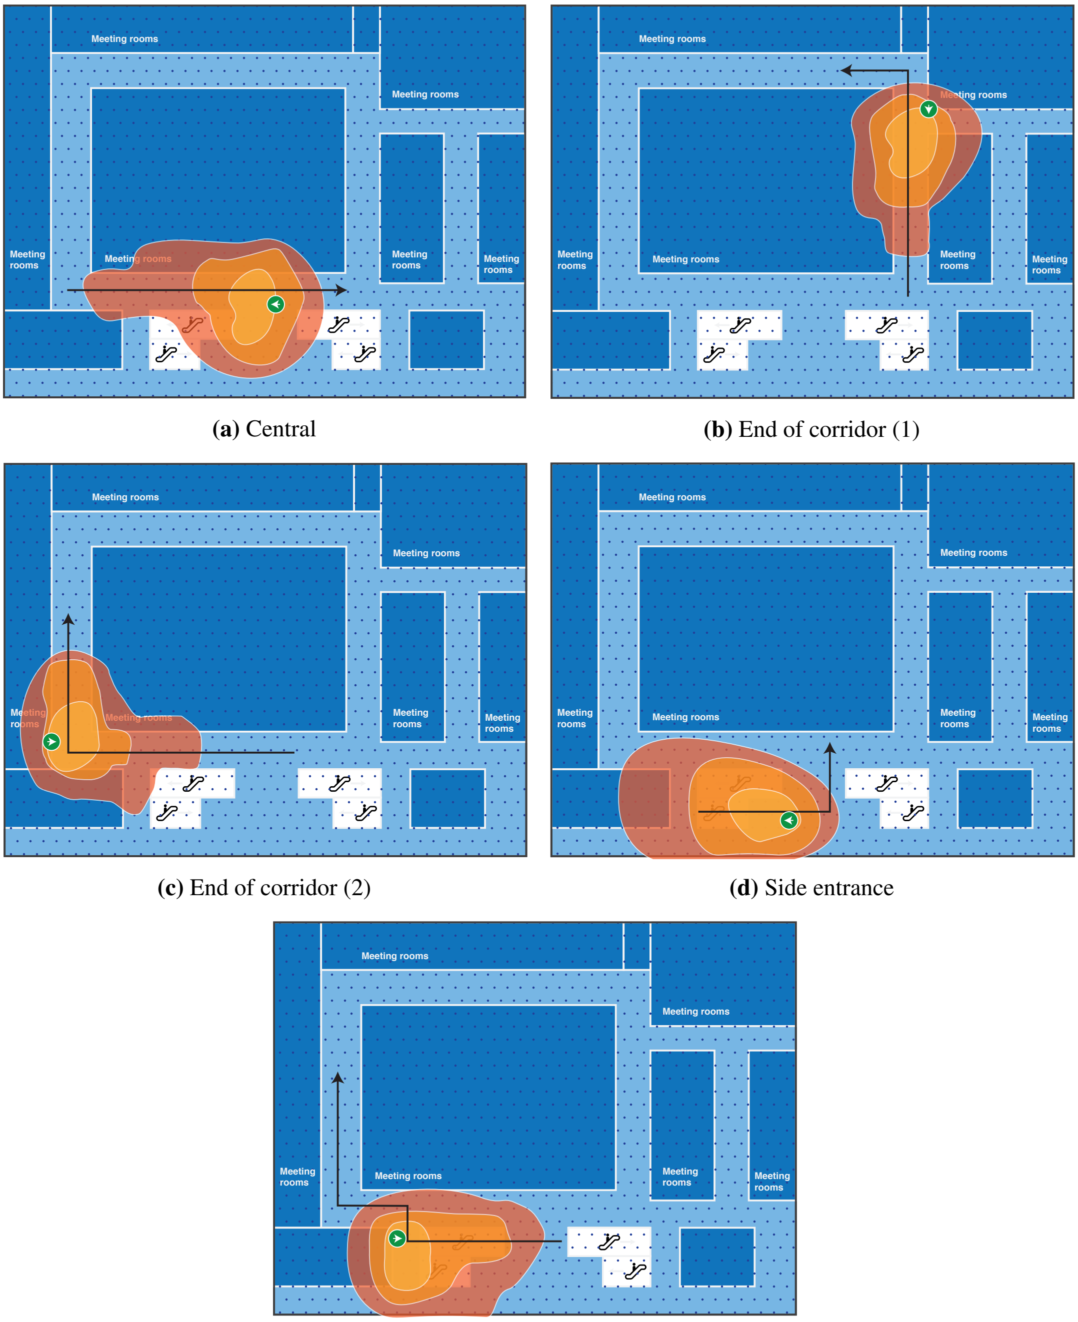

(a) Central

(b) End of corridor (1)

(c) End of corridor (2)

(d) Side entrance

(e) Main entrance

Figure 6.16: Floor plans showing the main floor of the convention space for each of the five experimental settings visualising the display location and its orientation (green circle), differently sized trigger zones (small, medium and large), and the paths of the walk-by experiments. Dots represent an approximation of the granularity of location points provided by the Wi-Fi location system (the floor plans provide an approximation only and have not been drawn to scale).

End of corridor (2) This display location is similar to 'end of corridor 1' but is located at the opposite end of the conference venue.

Side Entrance Whilst the previously mentioned display locations focused on the viewer remaining on the same level throughout the experiment, we further investigated the potential interference of the viewer changing floor levels throughout the experiment. This display is located on the main floor of the convention centre, however, the viewer appears in the trigger zone of the display (and enters the boundaries of the location tracking coverage) through an escalator from the level below.

Main Entrance To further investigate the potential interference of the viewer changing levels, we chose the displays at the main entrance as the final location. In contrast to 'side entrance escalator', the viewer transitions across three levels of which the bottom level and the main floor are covered by the location tracking system. The display is located on the main floor just at the end of the escalators.

In addition to different display locations and their characteristics, we varied the trigger zone (i.e. the spatial boundaries which viewers have to cross to be recognised in the viewable area of the display and trigger a content request) from small (i.e. trigger zone is identical to the viewable area of the display), medium (i.e. trigger zone approximately doubled in size) and large (i.e. trigger zone additionally increased and optimised for a high content exposure). For each display location and trigger zone size, we conducted ten repeating walk-bys, capturing the following timestamps: viewer entering the viewable area of the display, viewer leaving the viewable area of the display, content appearing on the display (i.e. the time the system recognised the viewer in the viewable area of the display) and content reverting back to the normal content. In the subsequent section, we provide detailed insights into the accuracy and performance of the system by focussing on content exposure and content accuracy metrics.

6.4.2 Impact of Trigger Zone Sizes

We begin our analysis by exploring the potential impact of different sized trigger zones on the performance of the overall system aggregated across all display locations and repetitions. Considering the box plots for each trigger zone type showing content accuracy and content exposure (Figure 6.17 and Table 6.6), we observe an increase of both content accuracy and content exposure with the size of the trigger zone. In particular, the small trigger zone (equalling the size of the viewable area of the display) shows very low content accuracy and content exposure measures (mean: 0.027 and 0.047 respectively), highlighting the limitations of the current location tracking technique. Medium- and large-sized trigger zones provide better results, whilst the large trigger zone reaches the highest means of 0.166 and 0.384 for content accuracy and content exposure respectively.

In ideal settings in which viewers are detected with the highest accuracy, we would expect an inverse correlation between the content accuracy and the size of the trigger zone (i.e. larger trigger zones leading to a decrease of the content accuracy). However, we observed a reverse

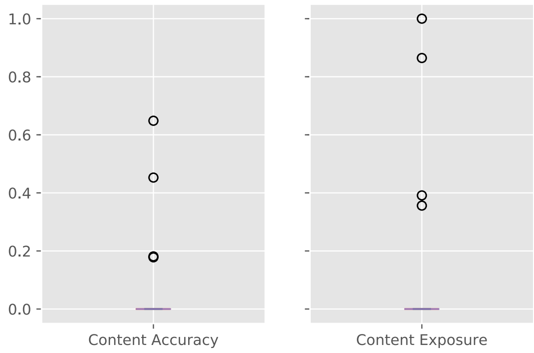

(a) Trigger zone: Small

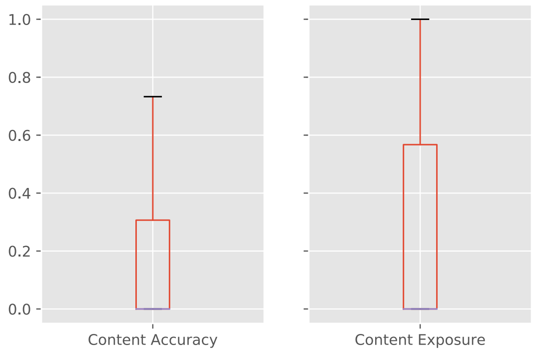

(b) Trigger zone: Medium.

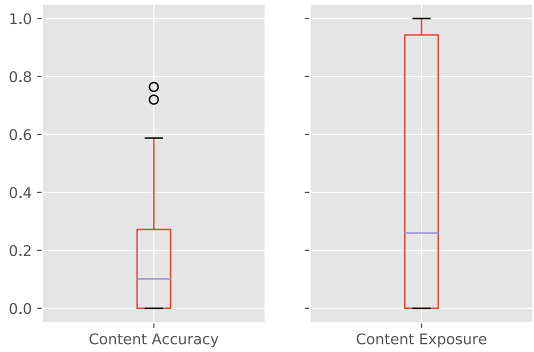

(c) Trigger zone: Large.

Figure 6.17: Box plots of accuracy and content exposure measures across all display locations for each trigger zone size.

Table 6.6: Content accuracy and content exposure measures across all display locations for each trigger zone size.

	Content Accuracy			**Content Exposure**		
Zone Size	Mean	Median	SD	Mean	Median	SD
Small	0.027	0.000	0.110	0.047	0.000	0.188
Medium	0.138	0.000	0.213	0.271	0.000	0.400
Large	0.166	0.101	0.203	0.384	0.260	0.417

effect to our assumption in which the content accuracy and content exposure improved with the size of the trigger zones (Table 6.6) – the low content accuracy measure likely caused by the delayed detection of the viewer before entering and after leaving the viewable area. Both the low content accuracy and content exposure metrics suggest that viewers can be detected soon enough to deliver content in time, however, the entry and exit detection latencies and accuracies appear to be low potentially leading to inaccurate analytics insights. In addition to the impact on accurately detecting a viewer approaching a display, the use of large trigger zones further impacts the potential scalability of the system. For example, if a number of displays are present in a constrained space, such displays may likely have overlapping trigger zones if the zones were configured to be very large. This consequently increases the number of potential viewers present in the viewable area of multiple displays, making it challenging to decide when to show which piece of content – and to capture individual display sightings and dwell times.

For digital signage analytics, it is particularly important to accurately detect when viewers enter the viewable area of a display – for example, to create accurate reports on audience sizes and the number of viewers passing by. In order to better understand the delay in which the underlying location tracking system detects viewers entering the space, we analysed the time from the viewer entering the viewable area to the time the system detects the viewers' location inside the trigger zone. Figure 6.19 shows a density plot of the delays for each trigger zone size – negative delays indicate that viewers were detected *before* entering the viewable area of the display while positive delays indicate that viewers were detected *after* entering the viewable area. In general, we observe an improvement of the detection latency for small (mean: 20.257, median: 20.4, SD: 14.828 seconds), medium (mean: 12.664, median: 8.7, SD: 20.895 seconds) and large trigger zones (mean: 6.781, median: 5.4, SD: 18.946 seconds). We note that with the increased trigger zone sizes, however, we observed a wider spread of the detection latencies indicating a high amount of noise and variance – leading to lower consistency in the detection of viewers in trigger zones.

6.4.3 Impact of Display Location Characteristics

To further investigate the potential impact of different display locations, we assume the 'best case' and only consider measures resulting from large trigger zones for each display location – allowing us to eliminate the potential impact of the accuracy and latency of the underlying location tracking system.

By using content exposure as the primary metric, we identify Main Entrance as the best performing display location (mean: 0.701; SD: 0.427). Whilst Central, End of Corridor (1) and End of Corridor (2) yield similar average performance results, we observe Side Entrance to be the display location with the lowest measures despite the large trigger zone (mean: 0.153; SD: 0.304). The low performance is likely a result of the unique characteristics of that location: viewers suddenly 'appear' in the viewable area of the display through an escalator that starts outside of the coverage area of the location tracking system. In contrast, whilst

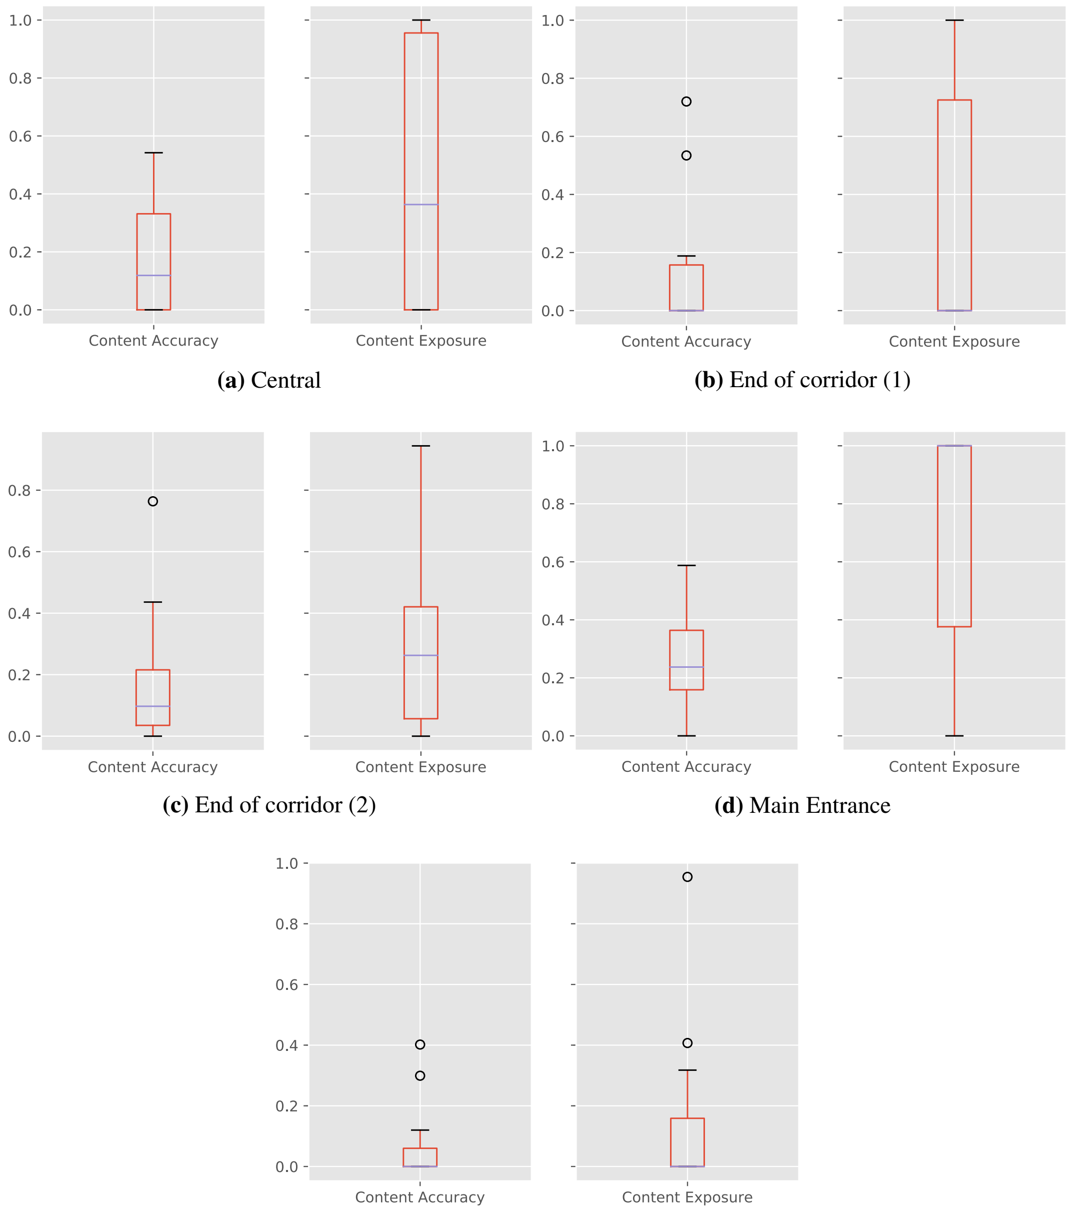

(a) Central **(b)** End of corridor (1)

(c) End of corridor (2) **(d)** Main Entrance

(e) Side Entrance

Figure 6.18: Box plots of content exposure and content accuracy measures for each display location and large trigger zone.

Table 6.7: Content accuracy and content exposure measures for each display location and large trigger zone.

	Content Accuracy			Content Exposure		
Location	Mean	Median	SD	Mean	Median	SD
Central	0.174	0.119	0.198	0.437	0.363	0.449
End of corridor 1	0.142	0.000	0.230	0.348	0.000	0.435
End of corridor 2	0.184	0.097	0.232	0.283	0.263	0.288
Side Entrance	0.075	0.000	0.143	0.153	0.000	0.304
Main Entrance	0.260	0.237	0.182	0.701	1.000	0.427

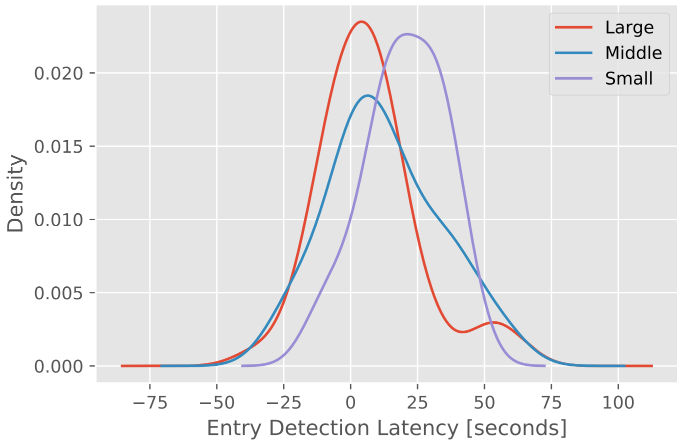

Figure 6.19: Density plots showing the delay of content turning on after/before the viewer enters the viewable area of the display for each trigger zone size across all display locations.

viewers also 'appear' in the viewable area of the display located at the Main Entrance, both the start of the escalator and the entire journey are within coverage of the location tracking system – leading to an improvement in the observed measures. For any location, the observed content accuracy appears to be rather low – likely an impact of the trigger zone that is significantly larger than the viewable area of the display. Additionally, we note that across all locations standard deviation measures were high suggesting a high variability in the captured measures.

To summarise, we observe that display locations and walk paths in which the viewer has already been present in the space and detected by the infrastructure-based location tracking system prior to approaching the display yield higher content exposures. We believe that this is a result of the limitations emerging from using Wi-Fi-fingerprinting and the low accuracy such systems currently impose. In order to counter balance the low accuracy, the use of large trigger zones may become necessary – further impacting on the usability of the system an metrics for analytics purposes and limiting the scalability of the system regarding large counts of displays and viewers.

6.5 Comparison of Viewer Mobility Tracking Approaches

As introduced in Section 3.4 (Capturing Viewer Mobility Data, p. 65) and evaluated in the previous sections, viewer- and infrastructure-based tracking are two approaches that can be utilised in order to capture mobility traces of viewers to provide a foundation for novel digital signage analytics. In this section, we describe the usability, benefits and costs of both approaches for each of the four stakeholder groups (display owner, space owner, content provider and viewer) and provide an overview of overall stakeholder independent benefits and drawbacks for each of the approaches.

6.5.1 Display Owners

Viewer-based Tracking

The instantiation of viewer-based tracking introduced in the context of this thesis utilised Bluetooth Low Energy beacons as the preferred location tracking mechanism. Of course, Tacita as the underpinning system is technology-agnostic and can be implemented using alternative location tracking mechanisms. If the location technology chosen requires the installation of appropriate sensors, costs for display owners may consist of the equipment and installation. Regardless of the technology chosen, display owners have to further provide the locations of their displays and associated capabilities (e.g. provision of personalised content) to the relevant viewers in the form of, for example, mobile phone applications or through other means.

The use of viewer-based tracking allows display owners to be independent from space owners as a distinct stakeholder group regarding the provision of appropriate location tracking technologies such as Wi-Fi fingerprinting. The use of viewer-tracking approaches can therefore be appropriate in environments where infrastructure-based tracking is not feasible (e.g. when appropriate Wi-Fi hotspots are not available), space owners are not willing to share location data due to commercial interests or legal reasons, or for very large display deployments that would require the integration of infrastructure-based tracking systems from a large number of independent spaces.

Infrastructure-based Tracking

The costs of using infrastructure-based tracking in the context of public display deployments lies particularly around access and integration. Display owners are required to integrate two potentially distinct systems (i.e. the display network with backends that provide location information of visitors present in a space) – an approach that does not scale with the number of distinct spaces in which displays are deployed.

The use of infrastructure-based tracking, however, allows display owners to draw on existing technologies and minimises the need for installations of additional sensing technology. Additionally, the use of mobile phone clients in order to enable viewers access to the system are not required in the examples provided in the context of this thesis. The use of infrastructure-based tracking is therefore appropriate in spaces that already provide all necessary capabilities, and where display owners wish to instantiate location tracking that is less reliant on viewers and their acceptance of potential third-party applications required in order to conduct the location tracking.

6.5.2 Space Owners

Viewer-based Tracking

The use of viewer-based tracking approaches may have impacts on the space owner. For example, the placement of displays equipped with sensing or beaconing technology in order to enable viewer-based tracking may raise some concern among both space owners and visitors of spaces. If viewer-based tracking is utilised in part in order to enable the provision of personalised content (such as via Tacita), space owners may further be impacted by the changing content displayed.

However, the use of viewer-based tracking may provide a cost benefit and shift of responsibilities to space owners. In order to achieve location tracking, space owners are not required to equip their spaces with appropriate tracking technology but can rely on other stakeholders. This is particularly beneficial if the installation of location tracking hardware within the space is technically, financially or legally not feasible.

Infrastructure-based Tracking

Drawing on the example of Wi-Fi fingerprinting as an approach to infrastructure-based location tracking of viewers, space owners are required to equip their space with an appropriate set of Wi-Fi base stations that are capable of supporting Wi-Fi fingerprinting. Additionally, space owners are required to provide access to interfaces of such base stations in order to enable the computation of locations of devices that are present in the space and connected to the Wi-Fi network. In order to ensure the accuracy and performance of the location tracking, regular tests and calibrations may be required.

The use of infrastructure-based tracking has a number of benefits. In particular, space owners are in full control and ownership of the location data captured. In contrast to viewer-based tracking, location data origins from a known and trusted source (e.g. systems provided by space owners or contracted third-parties). However, depending on the infrastructure-based technologies employed (e.g. Wi-Fi fingerprinting), visitors are still required to carry a mobile device in order to enable the location tracking within the space. The use of infrastructure-based tracking technologies can be the appropriate choice for space owners if the required hardware capabilities are available or can be installed.

6.5.3 Content Providers

Viewer-based Tracking

In the context of this thesis, content providers benefit from viewer mobility traces by enabling reports regarding content views across a display network, and enabling the provision of personalised content to viewers based on their preferences and locations. In order to utilise mobility traces, however, content providers are required to implement appropriate application programming interfaces. To enable the creation of viewer-centric reports (e.g. insights into the network visibility of content across a display deployment), content providers have to further supply access to logs of, for example, viewers or displays requesting certain types of content items.

By utilising viewer-based tracking to obtain access to mobility traces, the benefits for content providers lie particularly on the limited use of single or small number of client applications (e.g. Tacita) minimising implementation and data integration work.

Infrastructure-based Tracking

In contrast to viewer-based tracking, with the infrastructure-based tracking approach content providers are required to obtain mobility traces from individual space owners as the likely owner of the data produced. If content providers serve large numbers of distinct display deployments likely placed in environments that are owned and controlled by different stakeholders, the implementation and integration work is likely to grow with each deployment site. Different spaces are likely to heterogeneous location tracking technologies and interfaces for accessing the data.

However, obtaining mobility traces from space owners provides a trusted source of information. In particular example of the infrastructure-based tracking introduced in the context of this thesis, content providers are further able to access comprehensive location traces that go beyond of simple display proximity sightings of viewers as provided by Tacita. Such data can be utilised by content providers to better understand the potential impact of content shown to passers by (e.g. measure the impact of advertisement campaigns by analysing the viewers location traces after seeing a specific content item).

6.5.4 Viewers

Viewer-based Tracking

In order to support Tacita and enable the creation of meaningful insights, viewers are required to install a dedicated mobile phone application (e.g. the Tacita Mobile Client) and, in addition, allow the client application to access the viewers' location while the application is in the background (this is required in order to allow the application to listen for nearby beacons while the phone is in standby). The installation of a client application may have a number of implications for viewers. The background location tracking and processing has likely an impact on the battery lifetime due to additional energy use. Furthermore, the use of background location tracking can cause an invasion of the viewer's privacy.

By using viewer-based tracking, however, viewers are able to deactivate location tracking at any time. In the example of Tacita, viewers can choose between the deinstallation of Tacita from their mobile device, revoking of the permission to access their location, or by turning off location tracking within the client application itself. Furthermore, certain applications may provide a direct benefit to viewers. For example, Tacita allows viewers to request personalised and more relevant content when approaching displays nearby.

Infrastructure-based Tracking

In contrast to viewer-based tracking and the use of Tacita, infrastructure-based tracking takes away control from viewers and, in the context of Wi-Fi fingerprinting, it is sufficient for viewers to connect their mobile devices to the Wi-Fi network. Their location data is then automatically computed on the infrastructure side in real time. This can cause significant privacy-related issues where viewers may not necessarily be aware of the location tracking capabilities taking place in the environment. Due to the lack of awareness, it may be not obvious how viewers can (temporarily) opt out and deactivate location tracking in order to protect their privacy.

With the use of infrastructure-based tracking, however, viewers are also not required to install a dedicated client application reducing the burden of use. Furthermore, due to the data capture and processing taking place on the infrastructure, impacts on battery lifetimes and energy use are reduced to a minimum whilst providing a comparable service to the viewer.

6.6 Lottery Scheduling

The Lottery Scheduler was initially introduced in Section 5.3 (Lottery Scheduling for Digital Signage, p. 114) as a mechanism for supporting analytics-driven digital signs. We evaluate the Lottery Scheduler by investigating its applicability for typical digital signage networks and deploy it in the context of e-Campus. In particular, we analyse the ratio-based scheduling approach implemented in the form of a ratio-based lottery ticket allocation component. We focus on the performance of the lottery scheduler system and its content scheduling accuracy. The performance of the system is particularly important to support prompt responses and reactions to contextual changes that may yield immediate content changes such as content personalisation.

Excerpts of this section are based on the following publication:

- Mateusz Mikusz, Sarah Clinch, and Nigel Davies. "Are You Feeling Lucky?: Lottery-based Scheduling for Public Displays". In: *Proceedings of the 4th International Symposium on Pervasive Displays*. PerDis '15. Saarbruecken, Germany: ACM, 2015, pp. 123–129. ISBN: 978-1-4503-3608-6. DOI: 10.1145/2757710.2757721. URL: http://doi.acm.org/10.1145/2757710.2757721

6.6.1 Benchmarking

6.6.1.1 Apparatus

We performed our benchmarks on a typical e-Campus machine – a Mac Mini with the following specifications: 2.6 GHz Intel Core i5, 8 GB 1600 MHz LPDDR3 SDRAM; 1 TB HDD. The machine runs MacOS 10.10.2 (Yosemite) with Yarely as a digital signage player consisting of the Lottery Scheduler as the scheduling component. To understand the impact of the numbers of lottery tickets and content items, we performed a benchmark of the system with 1,000 and 10,000 tickets respectively and varying numbers of content items. We used simple dummy content items (images of 1000x1000 pixels and 0.6 MB in size) and repeated each combination of lottery tickets and content items 30 times.

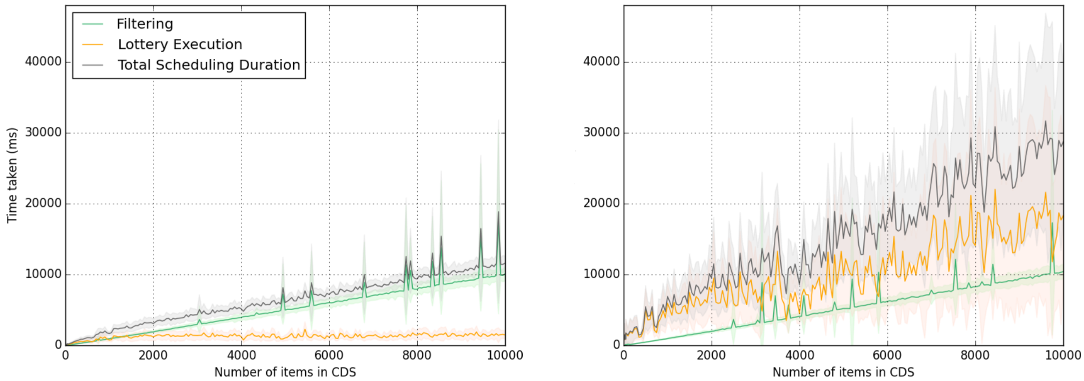

(a) 1,000 lottery tickets.

(b) 10,000 lottery tickets.

Figure 6.20: Lottery scheduler lab-based benchmarks with 0-10,000 content items with low and high numbers of lottery tickets (initially published in [MCD15]).

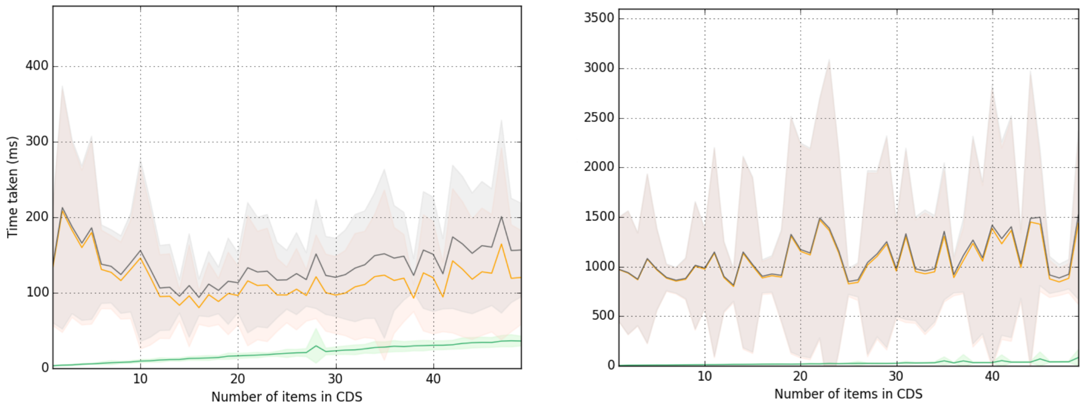

(a) 1,000 lottery tickets.

(b) 10,000 lottery tickets.

Figure 6.21: Lottery scheduler lab-based benchmarks with 0-50 content items with low and high numbers of lottery tickets (initially published in [MCD15]).

6.6.1.2 Results

The results of the performance benchmarks are visualised in Figure 6.20. Overall, an increased amount of lottery tickets available to allocate to content items leads to an increased delay in the time taken to perform the lottery ticket allocation and to yield a content item to schedule. The time taken to finish the ticket allocation process increases by a mean of 7.77 with a ten-fold increase of lottery tickets (Figure 6.21). This increase is likely a result of our implementation choice for the lottery scheduler: the ticket allocation process utilises Python-based Queue objects. Therefore, each lottery ticket allocation results in a context switch across Python threads performing the ticket allocation and provides an overhead in the processing time. The filtering of ineligible content items due to constraints defined in the Content Descriptor Set is performed prior to the ticket allocation and is therefore independent of the number of available and allocated lottery tickets.

In more detail, the increase in the number of content items is directly correlated with the time taken to complete content scheduling decisions.

The mean time taken of the lottery (including ticket allocation and random draw) for $1,000$ tickets are 181.60, 736.0, 1477.47 milliseconds for 100, 1000 and 10000 content items. Figure 6.20a visualises the linear relationship between the number of content items and the duration of the total scheduling decision mainly dominated by the filtering time. Considering the time for completing the lottery execution, we observe that at the point at which the number of content items outgrows the number of available lottery tickets the time to execute the lottery duration remains constant (Figure 6.21a). In the case in which the number of content items is smaller or equal to the number of available content items, we observe the time taken to complete the lottery execution to be directly correlated with the number of lottery tickets (Figure 6.21b).

Considering the filtering and other components of Yarely in the benchmarking, we see that the performance of other Yarely components is unaffected by lottery scheduler components with the increased number of lottery tickets. However, some components are affected by the increased number of scheduled content items. The filtering of ineligible content items takes 171.47, 950.57 and 10405.70 milliseconds for 100, 1000, and 10000 content items respectively. This is due to the implementation of the filtering component that loops through every single scheduled content item and applies a range of filters. However, as shown in Figure 6.20 the time complexity increases linearly with the number of content items. In previous work, the overall time taken to complete a scheduling decision in Yarely was measured with approximately 1.5 seconds (including a 0.6 second fade animation between content items) [Cli+13].

The overall performance of the content scheduling process (including filtering, lottery ticket allocation, and randomised draw of the 'winning' content item) is significantly faster for lower numbers of content items. For 100 content items, the overall scheduling process is completed in less than 0.3 seconds using $1,000$ tickets. In contrast, larger numbers of scheduled content items take noticeably longer to complete the scheduling process. The overall scheduling time for $1,000$ and $10,000$ content items with $1,000$ tickets each are 1.66 and 11.58 seconds respectively. In our experience from the e-Campus display network, however, we found that displays are typically subscribed to approximately 30 content items each (mean 31.33, median 28.5, max. 107). The lottery scheduler therefore meets the performance expectations of typical display networks. Of course, the performance for larger amounts of content items can be improved in various ways. For example, the filtering and lottery ticket allocation processes can be initiated ahead of time instead of at the time at which the content decision is required, i.e. well before the currently played content item reaches the end of its content playback time. However, the drawbacks can include a reduced ability to react to changes in the context of the display that may take place before the new content decision is required but after the scheduling of the subsequent content item has been initiated.

6.6.2 Accuracy in a Real-World Deployment

6.6.2.1 Apparatus

To further investigate the accuracy and reliability of content scheduling requirements (Section 5.3.1, p. 114), we investigated the performance of the Lottery Scheduler in a real-world setting in the context of e-Campus. The content and display management system for e-Campus features the ability to group individual content items into 'Channels'. Additionally, display owners can specify content playback ratios per channel if multiple channels have been scheduled onto a single display to prioritise individual channels over others (described in further detail in Section 1.3, p. 6).

Utilising the channel system and existing content, we created four content channels within the Channel system and allocated a demonstrative set of content items from content available in the Channel system. In particular, the four channels were configured with the following parameters:

- **Channel A:** 16 images, 0.51 playback ratio / air time,
- **Channel B:** 4 images, 0.25 playback ratio / air time,
- **Channel C:** single 60-second video, 0.12 playback ratio / air time, and
- **Channel D:** 4 images, 0.12 playback ratio / air time.

We configured the Lottery Scheduler with a single lottery ticket allocation module for the ratio-based allocation of lottery tickets, and provided $10,000$ tickets for allocation. The scheduler was executed on a single display for continuous 12 hours subscribed to the channels described above. With regards to the implementation of the Lottery Scheduler, all content items were played for their full duration, i.e. the video for its full duration of 60 seconds and all other content item types for a default content playback time of 15 seconds. The total duration of all content items across the four channels was therefore 420 seconds. In the analysis, we consider the accuracy of the content playback with regards to predefined ratios.

6.6.2.2 Results

We first consider the playback ratios produced by the ratio-based ticket allocation approach in comparison with the expected (i.e. pre-configured) channel ratios described in the previous subsection. As shown in Figure 6.22, the overall channel ratio accuracies are close to the expected ratios after an initial phase of high variations. In particular, the system approximates expected ratios (0.51; 0.25; 0.12; 0.12 for Channels A, B, C and D respectively) within a short time period:

- **After 15 minutes:** 0.65; 0.16; 0.07; 0.12,
- **After 30 minutes:** 0.59; 0.22; 0.11; 0.08,

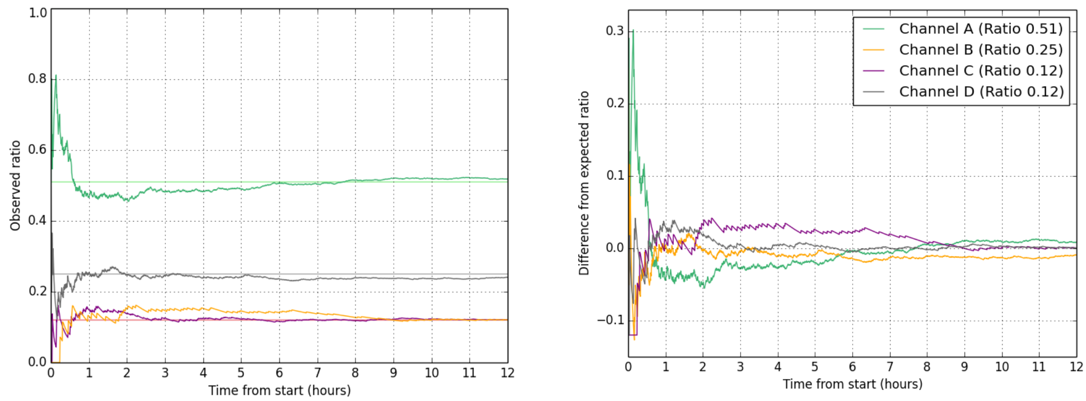

(a) Absolute ratios per channel – paler marks indicate expected ratios.

(b) Deviation from observed and expected ratios per channel.

Figure 6.22: Accuracy of the lottery scheduler using a ratio allocator and sample e-Channel content (initially published in [MCD15]).

- **After 45 minutes:** 0.50; 0.25; 0.13; 0.13, and
- **After 720 minutes (12 hours):** 0.52; 0.24; 0.12; 0.12.

Over time, the observed ratios converge toward the predefined ratios: for example, the discrepancy between the expected and observed ratios for channels A, B, C, and D was measured after 45 minutes at 0.01, 0.00, 0.01 and 0.01 respectively. The lottery scheduler requires a short initialisation period only (approximately 15-30 minutes) to provide an acceptable level of accuracy – we note, however, that the random draw of content items across channels has a small impact on the observed ratios.

Considering the playback of individual content items within channels, we observe that each content item is played almost equally. This reflects the implementation of our ratio-based ticket allocation module: content items within an individual channel receive an equal amount of lottery tickets (based on their content playback length) and have therefore equal probabilities in the random draw with regards to the overall playback time. As an example case, we further investigated the content playback times within Channel D. Overall, we observe that content items from Channel D have only occupied around nine minutes of content playback time within the *first hour* of the accuracy benchmarking whilst the four items within Channel D have occupied between 1.16 and 3.23 minutes, i.e. between 12.86 and 35.77% of content playback time of the channel – in contrast to the expected content playback times as an even distribution of 2.25 minutes each, i.e. 25%. After *four hours*, the playback times of content items within Channel D have further improved. The channel occupied approximately 30 minutes of the total four hours of content playback, whilst each individual item within the channel played for 6.00-8.91 minutes (i.e. 20.04-29.79% of the channel playback time). Towards the end of the twelve hour experiment, Channel D has occupied approximately 87 minutes whilst each item consisted of a playback time between 16.16 and 25.70 minutes (i.e. 18.64–29.65% of the channel playback time).

In addition to the accuracies with respect to content playback ratios, we were able to observe behaviour not typically seen in other content scheduling approaches such as round robin. Due to the nature of the Lottery Scheduler of using a random draw to determine the 'winning' content item for playing on the display, in some cases the same content item is shown on the screen consecutively. Within our 12 hour experiment, we were able to observe content items to be played for up to three times in a row, resulting in an increase of the playback time of an individual image from 15 to 45 seconds in our configuration. We note that such behaviour can be overcome by utilising additional filtering or ticket allocation modules. For example, a filtering component could consider a previously played content item in the current round of content scheduling as an ineligible content item and exclude it from content scheduling until a subsequent round. Alternatively, an additional ticket allocation module can be implemented that prioritises content items that have not been played within a certain time period. However, such filtering and ticket allocation modules in combination with the ratio-based ticket allocation may lead to a higher discrepancy between the predefined and observed channel ratios.

6.7 Summary

In this chapter, we described a set of trials that sought to provide real-world insights into the effectiveness of our display analytics technology probes. In particular, we have made the following set of contributions.

- We conducted an evaluation of Pheme in which we presented the integration of Pheme and its client libraries into the e-Campus display test-bed, and collected a large quantity of data over a period of over four years. We showed the feasibility of the Pheme architecture for leveraging existing third-party analytics engines by providing insights into the implementation of the Google Analytics injector and its deployment in the context of e-Campus.

- We conducted a thorough evaluation of collecting viewer-centric analytics data and providing personalised services to individuals in the context of Tacita. We highlighted the benefits and limitations of the system by investigating the accuracy and performance in the context of an in-the-wild deployment, and particularly focussed on systems-related benchmarks and viewer detection latencies. We showed that it is feasible to use the system to capture accurate insights about individuals, and fast enough to support walk-by personalisation.

- We further evaluated Wi-Fi fingerprinting in the context of a large, commercial convention centre as an alternative mechanism to capturing viewer behaviour and movement patterns. We conducted a set of walk-by experiments in the context of a large convention centre providing insights into the performance and accuracy of the underlying location tracking system and the overall system implementation.

- We evaluated the performance of the lottery scheduling approach through a set of benchmarks, and evaluated the system through a long-term and in-the-wild deployment in the context of the e-Campus display deployment. We demonstrated the feasibility and reliability of the lottery scheduling approach as the main scheduler for e-Campus.

In the next chapter, we provide an overview of the contributions of the dissertation as a whole, and elaborate on future directions and opportunities in the areas of digital signage analytics.

Chapter 7

Analysis, Conclusions and Future Work

7.1 Overview

In Chapter 1, we described the emergence of digital signage and pervasive display networks and the importance of gaining detailed insights into the impact of digital signs and content on viewers. Whilst in related domains such as Web analytics it has been possible to measure the cause and effect of, for example, advertisement campaigns through comprehensive tracking and analytics mechanisms, in the physical world it is less obvious how the direct cause and effect of displays on individuals can be measured. Motivated by a scenario of a local shop owner which identified the potential of future digital sign analytics that could capture the cause (viewer seeing a piece of content on a display) and its effect (offline purchase of the advertised product), we described the vision of viewer-centric analytics.

In Chapter 2, we provided an overview of related work in the areas of digital signage analytics. We first introduced audience models and metrics that have been developed to describe viewer interaction in front of individual displays such as the Audience Funnel developed by Michelis and Müller [MM11]. We then provided an overview of data capture techniques (e.g. Intel AVA [Cav11] that utilised video computing technology to capture audience numbers and demographics), ways to report analytics data (e.g. funnel and flow diagrams), systems and deployments that utilise analytics data for display actuation purposes (e.g. targeted advertising) and the use of digital signs in specific contexts (such as retail). The majority of previous work focusses on individual displays and are limited in their ability to capture overarching analytics and provide insights into the impact of digital signs and the content shown.

Looking to explore analytics beyond individual signs, in Chapter 3 we conducted an extensive literature review identifying categories of data that can be captured by different stakeholders in open pervasive display networks. We provided a framework that supports the description of potential combinations of analytics data owned by distinct stakeholders. Utilising the framework, we highlighted the potential benefits that emerge from sharing analytics data across these stakeholders and identified the importance of new viewer-centric analytics in the context of open pervasive display networks. We provided insights into the design and development of Pheme, an analytics platform that supports leveraging third-party analytics systems through the specification of injection modules. We also presented the design and development of two systems enabling the capture of viewer mobility data: a viewer-centric approach (including a corresponding back-end system design and implementation) in which we utilise client-based tracking technology captured through viewer mobile phones, and an infrastructure-centric approach in which we designed a middleware platform that utilises viewer mobility tracking through Wi-Fi fingerprinting. As both approaches require the tracking of individuals within a space and potentially introduce a number of privacy-concerns, we introduced synthetic analytics as a privacy-friendly alternative. This approach uses mobility models to create synthetic viewer mobility traces and combines these with real-world analytics events.

In Chapter 4, we focussed on the creation of new analytics reports for open pervasive display networks. We began by presenting a set of novel viewer-centric analytics reports that illustrated the levels of insight that can be gained when utilising viewer-sightings (captured through either viewer-based tracking or mobility models). Examples of such reports provide insights regarding the effectiveness of displays (e.g. the number of unique viewers per display), the visibility of content across the display network (e.g. the number of unique viewers per content item across the network), and the visibility of content to viewers (e.g. number of times the same content item was seen by viewers across the network). Drawing on our work on capturing viewer mobility traces, we describe a set of new analytics reports specifically relating to display personalisation systems. Finally, we illustrated how existing Web analytics engines can be leveraged (using appropriate Pheme injection modules) to create a new set of display-oriented analytics reports. We presented a mapping from display to Web analytics terminology, and designed and developed an appropriate injection module for Pheme that implements the mapping. We provided example reports through Web analytics regarding the displays in the network and content shown.

Whilst many analytics systems focus on the creation of analytics reports as the end product, we continued our exploration in Chapter 5 into the automated use of analytics data on digital signs and designed and developed systems to support the use analytics data to, for example, drive decisions for content scheduling. Our Lottery Scheduler provides a novel content scheduling approach with the ability to support context- and event-based scheduling allowing the display to instantly react to contextual changes in the environment.

In Chapter 6, we provided evidence of the feasibility of the approaches and systems introduced in previous chapters spanning across *data capture*, *reporting* and *automated use of analytics data*. We began the evaluation with Pheme, our analytics support platform, by providing insights into its integration into the e-Campus displays network, the real-time mapping and injection of events into third-party analytics engines, and the large quantity of data collected and processed (over 73 GB in textual data). We continued the chapter by providing evidence for the feasibility and accuracy of the Lottery Scheduler, and showed the

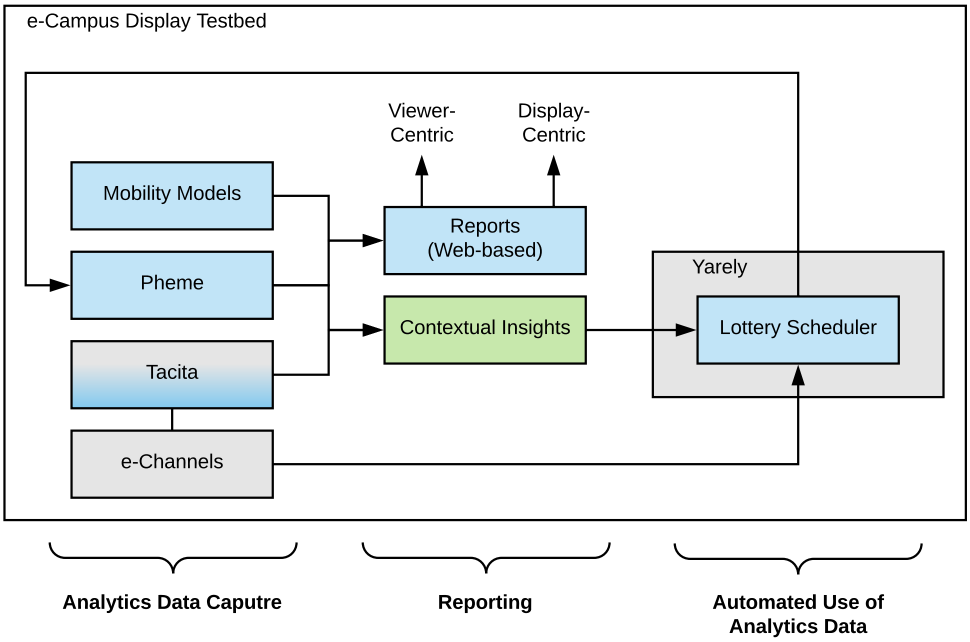

Figure 7.1: End-to-end implementation and integration of the systems introduced in this thesis. Blue: fully implemented components as part of this thesis; Green: partially implemented component; Grey: components and systems not developed in the context of this thesis but part of the e-Campus display network.

performance of the underlying system – crucial to support prompt content scheduling due to contextual changes. We continued to describe the design and implementation of the Tacita trial, an in-the-wild deployment of the display personalisation system as part of the e-Campus display deployment. We provided detailed insights into the performance of individual system components, the feasibility of using Bluetooth Low Energy sensing technology to detect and capture viewer analytics data (e.g. dwell times), and the overall use of display personalisation over an extended period of times (over 200 days and 150 users). We concluded the evaluation by investigating the feasibility of using infrastructure-based tracking as alternative to the Tacita approach by conducting a controlled in-the-wild deployment in the context of a large convention space. As part of this trial, we conducted a series of walk-by experiments providing insights into the overall accuracy and performance of the system, both relating to display analytics and to support display personalisation.

7.2 Analysis

7.2.1 Findings

In this thesis, we have explored techniques for the next generation of digital signage analytics systems. As part of this exploration, we designed, developed and deployed a range of systems and components and evaluated these, primarily in the context of the e-Campus display test-bed.

We did not aim to create an integrated digital signage analytics solution. Nevertheless, as shown in Figure 7.1, the majority of the systems and components that emerged from this thesis can be integrated into an end-to-end system. In particular, we have integrated Pheme, Tacita, Mobility Models and the Lottery Scheduler to create an analytics solution for e-Campus.

Pheme serves as the main platform and repository for capturing, processing and storing display analytics data. Individual display nodes report each displayed content item at the point at which the content was made available (via the Lottery Scheduler), changes in the physical state of the display and potential contextual changes in real time through the integration of Pheme client libraries into the relevant components of the Lottery Scheduler (described in detail in Section 6.2.1). Pheme processes and stores this data and, additionally, injects relevant portions of the incoming data stream into Google Analytics enabling the real-time report generation (in accordance to the mapping introduced in Section 4.3). These display-oriented reports are accessed by display owners, content providers and administrators for monitoring purposes. Additionally, the datasets stored in Pheme are combined with data captured through Tacita (Section 3.4.1) and Mobility Models through the Synthetic Analytics approach (Section 3.4.3) to create viewer-centric reports. Furthermore, Tacita pushes contextual events (e.g. content schedule requests) back into the set of displays. We note, however, that the components forming Contextual Insights (Figure 7.1, highlighted in green) have been only partially implemented and are currently limited to reporting display sightings of viewers to enable walk-by display personalisation (Section 6.3.1). Both content scheduling decisions and reported contextual changes on the Lottery Scheduler are reported back into Pheme in real time – 'closing the loop' of our digital signage analytics platform.

As part of our exploration into the collection and processing of relevant analytics, we highlighted the opportunities for collecting a range of analytical insights in the context of open pervasive display networks (Section 3.2). In particular, the large number of stakeholders of a display network and their access to a unique set of analytics data enable the computation of novel insights. We identified such opportunities both from our experiences in developing novel analytics for the e-Campus display network and a comprehensive literature survey (Section 3.2.2).

Finding. *A wide range of analytics data can be captured by stakeholders in an open pervasive display network. Synthesising such datasets across distinct stakeholders enables the creation of novel and comprehensive analytics reports allowing us to describe the effectiveness of displays, and providing insights into how viewers experience the display network as a whole. These novel analytics reports and the synthesis across stakeholders will be crucial to the success of future open pervasive display networks.*

We showed that both Web and digital signage analytics domain share common characteristics in the types of desired reports and levels of insights (Section 4.3.3 and Section 6.2). In this context, we explored the benefits gained for creating relevant analytics reports by leveraging existing Web analytics engines. To provide a proof of concept, we developed a mapping and corresponding Pheme injection module allowing us to use Google Analytics for the creation of display-oriented reports (Section 4.3.3).

Finding. *Digital signage analytics can benefit from the wealth of existing Web analytics engines. The opportunities from mapping display to Web analytics terminology have been highlighted through the creation of novel display-oriented reports within Google Analytics.*

By integrating and deploying Pheme to capture analytics events across the e-Campus display test-bed we were able to provide an overview of the current state of the entire display network in real time (e.g. by accessing real-time reports via Google Analytics). Such reports were particularly valued by stakeholders of the display network (both content providers as well as administrators).

Finding. *Pheme plays an important part in the daily maintenance of the e-Campus display network and is used to monitor displays in real time using Pheme's APIs and appropriate analytics reports.*

We further explored the creation and use of analytics reports that go beyond display-oriented analytics. In particular, we explored the capture and creation of viewer-centric analytics which consist of further insights about individuals such as their mobility patterns (Section 3.2, p. 49). We identified two categories of analytics datasets: display-oriented analytics data (originating from the signs themselves including the Lottery Scheduler), and viewer-oriented mobility patterns – forming the foundation for our novel viewer-centric analytics reports. In order to create a novel set of viewer-centric analytics reports, we highlighted the opportunities that emerge when traditional display-oriented analytics data from Pheme is combined with mobility patterns of viewers.

Finding. *The collection of traditional display-oriented analytics is an important prerequisite for the creation of novel, viewer-centric analytics reports that combine both display- and viewer-oriented data.*

Such mobility patterns can originate from various data sources (Section 3.4, p. 65). In particular, we utilised three distinct mechanisms to capture viewer mobility data: viewer-based tracking (e.g. through the viewers' mobile phones), infrastructure-based tracking (e.g. using Wi-Fi fingerprinting) and synthetic mobility traces (e.g. through appropriate mobility models).

The combination of both display- and viewer-oriented datasets enable us to gain insights into how viewers experience a digital signage network as a whole.

Finding. *By capturing and processing both display- and viewer-oriented analytics data, we can create novel reports regarding the effectiveness of displays, the visibility of content across the display network, and the visibility of content to viewers.*

We particularly emphasise the importance of the shift of perspective from display- to viewer-oriented analytics. In contrast to focussing on individual displays (e.g. capturing audience numbers and demographics for individual displays) and providing aggregates over all displays, we instead shift the focus to the viewer. How do viewers experience a display network and content as the navigate through and across spaces?

Finding. *Novel viewer-centric analytics for open pervasive display networks help us understand* ***how viewers experience the digital signage network****. For example, the order in which they see content displayed across the network, and the amount of times viewers see content repeatedly.*

In the context of e-Campus, we explored two distinct approaches: viewer-based tracking of display sightings through Tacita (Section 3.4.1) and Mobility Models used to support our synthetic analytics approach (Section 3.4.3). The utilisation of both approaches allowed us to create a set of viewer-centric analytics reports by additionally accessing datasets captured through Pheme (Section 4.2).

Finding. *The combination of capturing data on the viewer-side (e.g. mobility traces) and server-side (e.g. content logs) are sufficient to create novel and insightful analytics reports regarding the effectiveness and success of a digital signage network.*

Capturing viewer mobility traces, however, is highly technology-dependant and we explored a range of techniques as mentioned previously. By comparing these approaches, we highlight that the accuracy of analytics reports created based on synthetic viewer traces is highly dependant on the quality and accuracy of the underlying mobility models. Regarding the accuracy, provide a higher accuracy due to the underlying detection technique through BLE beacons – however, impose privacy concerns to the viewer.

Finding. *Capturing accurate viewer mobility traces can be best achieved through client-side tracking. However, the synthetic analytics approach provides a privacy-preserving alternative to capturing mobility traces and is suitable for environments in which fine-grained location tracking is infeasible.*

As part of our trails regarding the use of both viewer- and infrastructure-based tracking, we were able to capture performance measures for both approaches (Sections 6.3, p. 130; and 6.4, p. 149). Whilst the subsequent evaluations were not designed to explicitly compare both approaches (due to the differing evaluation locations for each of the tracking technologies), the results give us an insight into the different characteristics of both approaches.

Finding. *The collection of viewer mobility data through viewer-based tracking technologies (i.e. Bluetooth Low Energy beacons) provides the best accuracy (regarding the tracking of viewers) and performance (regarding the delay of detecting viewers in proximity to displays) compared to infrastructure-based tracking via Wi-Fi fingerprinting.*

A portion of our work focussed on exploring alternative uses of analytics insights beyond the creation of reports. The large amount of stakeholders contributing to open pervasive display networks, and, as a consequence, the amount of contextual events, requirements and constraints that need to be considered when creating content scheduling decisions lead to the need for novel content scheduling approaches that are capable of accommodating these constraints and requirements. In this thesis, we presented a scheduling approach for digital signage that uses Lottery Scheduling to address these requirements. The Lottery Scheduler was integrated as part of the e-Campus display test-bed and provides interfaces for both the e-Channel system (to request its regular content schedules) and Tacita to capture contextual events (Section 6.2); particularly display sightings of viewers who requested personalised content.

Finding. *The Lottery Scheduler provides a suitable mechanism to facilitate a range of (potentially competing) content scheduling requirements and constraints that can originate from a number of stakeholders whilst still dynamically responding to contextual changes and real-time analytics that influence content scheduling decisions.*

7.2.2 Benefits to Other Research Communities

A number of findings and insights that emerged from this work have potential benefits for areas of research besides distributed systems and pervasive computing.

Human-Computer Interaction In the context of this thesis we created a number of novel analytics technologies that could be utilised by the Human-Computer Interaction community to help understand user behaviour and engagement patterns of interactive systems beyond pervasive displays. For example, our approaches could be used to understand how users interact with future smart IoT environments without requiring extensive additional instrumentation or observation of these spaces.

Modelling User Behaviour As part of our large-scale in-the-wild trials, particularly Tacita, we were able to capture large amounts of viewer location traces from our deployment. These location traces can serve as a foundation for the creation of more realistic models of human mobility – e.g. to better understand and simulate human movements in the context of our university campus and informing future work. In this context, the data can be used to both train novel models and provide test data to evaluate the accuracy of the models developed and to provide feedback for potential machine learning approaches.

Mobile Computing Mobile computing researchers could benefit from the insights gained as part of this work to better understand how mobile devices can be used to support interactions with smart environments in general. Additionally, detailed insights into the temporal and spatial accuracy of Bluetooth Low Energy beacons for capturing user locations can inform future developments of location tracking technology and associated applications.

7.3 Contributions

In the context of this dissertation, we have made contributions in four key areas: 1. the collection of analytics-related datasets for digital signage systems, 2. the utilisation of collected datasets for the creation of analytics reports and the exploration of novel, viewer-centric analytics reports, 3. the development of a new scheduling systems supporting dynamic content scheduling based on contextual and analytics events, and, 4. a set of large scale and long-term trials in the wild investigating the technical feasibility and applicability of our systems.

In the subsequent sections, we summarise the key contributions in greater detail in the order in which the contributions appeared in the dissertation.

7.3.1 C1: Analytics Data Collection

The first set of contributions of this thesis relates to the collection and processing of analytics data that is fundamental for the creation of novel insights in the digital signage domain (Chapter 3, p. 48). This included, in particular, the following set of contributions.

1. We introduced a framework for the categorisation of relevant analytics data for the digital signage domain (Sections 3.2 and 3.2.3, p. 49 and p. 53 respectively). This included a detailed literature analysis to capture the categories of data that individual stakeholders of future pervasive display networks are able to capture (Section 3.2.2, p. 51). Based on the literature review and subsequent analysis, we further contribute the identification of opportunities that arise from synthesising analytics data from distinct stakeholder groups for the generation of novel and viewer-centric analytics insights (Section 3.2.4, p. 56).

2. We highlighted the importance of traditional display analytics as a foundation for creating novel analytics insights, and developed a set of client libraries that can be integrated into existing digital signage deployments to enable the collection of relevant datasets in real time (Section 3.3.2, p. 60). We designed a back-end systems architecture for the collection, processing, reporting and export of relevant datasets – including a proposed extension for future data synthesis (Section 3.3.3, p. 62).

3. A crucial step towards the creation of novel viewer-centric analytics is the collection of viewer mobility and behaviour data (Section 3.4, p. 65). Our initial exploration of the collection of viewer mobility data was conducted in the context of a display personalisation system that relies on viewer-side tracking through a dedicated mobile phone application detecting and reporting near-by displays through Bluetooth Low Energy beacons (Section 3.4.1, p. 66). We developed and deployed the system in the context of the e-Campus display network at Lancaster University.

4. We additionally designed and developed a system that supports the collection of viewer mobility data through infrastructure-based tracking mechanisms (Section 3.4.2, p. 73). The proposed middleware system was designed to integrate with existing tracking and display infrastructures, and was developed in the context of a large convention centre in Singapore.

5. Due to the potential technical limitations in tracking viewer mobility through both client- and infrastructure-based sensing technology, and the existing privacy concerns when tracking individuals within and across spaces, we introduced synthetic analytics as an alternative approach to the collection of viewer mobility data (Section 3.4.3, p. 76). We developed a system that transformed a spatial map into a graph-based data structure, and used a set of example mobility models to show the applicability of the synthetic analytics approach.

7.3.2 C2: Reporting

For our second set of contributions we introduced a set of novel viewer-centric analytics reports that build on top of the datasets from C1 (Chapter 4, p. 82). In particular:

1. We identified and created a set of novel viewer-centric analytics reports that provide new perspectives on the effectiveness of displays, the visibility of content across a display network, and the visibility of content towards individual viewers (Section 4.2, p. 83). These reports are primarily founded on the data captured through the synthetic analytics approach combined with display-related analytics from Pheme.

2. We describe a new set of analytics reports for novel signage networks that support the delivery of personalised content to viewers (Section 4.2.5, p. 98). Such reports include insights into usage and interactions, and the meaning of ‘usage retention’ in the context of digital signage deployments.

3. We highlighted that the type of analytics for digital signage closely mirrors the type of analytics reports that can be created in common Web analytics systems. In this context, we introduced novel types of digital signage analytics reports created by leveraging existing Web analytics engines. In particular, we provided a mapping from digital signage to Web analytics terminology and presented the implementation of an appropriate injection module for Pheme serving as a proof-of-concept that can feed signage analytics events into Google Analytics (Section 4.3, p. 102).

7.3.3 C3: Automated Use of Analytics Data

Our third set of contributions consists of the design and development of a novel content scheduling system for digital signage that lays the foundation for the automated use of analytics data such as the datasets introduced in C1 (Chapter 5, p. 111). In particular:

1. We highlighted the need for novel digital signage scheduling systems that support the automated use of analytics data in the context of open display networks. In particular, we emphasised the need for such systems to respond to a large number of potentially conflicting scheduling constraints and requirements that, for example, originate from the large set of stakeholders and contextual events taking place in the vicinity of the display (Section 5.2, p. 111).

2. We provided a description of the design and development of the first lottery scheduling system for digital signs (Sections 5.3, p. 114; and 5.4, p. 117). The Lottery Scheduler was specifically designed as a technique that allows the resolution of potentially conflicting scheduling constraints and requirements, and responds to dynamic contextual changes such as display sightings from Tacita (Section 5.5, p. 121).

7.3.4 C4: Systems Evaluation and Large-scale Trials

Our final set of contributions specifically focus on large-scale and long-term system evaluations conducted in the context of the e-Campus display test-bed (Chapter 6, p. 124). In particular:

1. We integrated the analytics data collection systems part of C1.2 into the entire e-Campus display deployment consisting of over 65 displays and captured over $159,264,530$ display-related events yielding a volume of 73.67 GB (excluding database indices) in the course of the trial (Section 6.2, p. 125). The deployment has been ongoing since August 2014 and serves as the main analytics platform for e-Campus demonstrating the viability of analytics for technologies presented in this thesis.

2. We conducted a large-scale and in-the-wild trial of a display personalisation system utilising client-side tracking technology described in C1.3 and C3.2 (Section 6.3, p. 130). Our contributions include in particular a detailed investigation into the feasibility of Bluetooth Low Energy beacon technology for capturing relevant and accurate viewer-related analytics insights and walk-by personalisation (Section 6.3.2, p. 134) as well as

general long-term evaluation of the overall system including middleware components and extensions to the existing display infrastructure (Section 6.3.4, p. 144). The system has been deployed for over 206 consecutive days, attracted over 147 users and yielded a total of $226,620$ events of which $24,565$ consisted of content request and an equal number of display sightings providing us with initial insights into viewer mobility and behaviour across the University campus.

3. As a comparison to client-side tracking, we additionally evaluated the system performance and feasibility of using infrastructure-based tracking technology that utilises Wi-Fi fingerprinting described as C1.4 and C3.2 in the context of a large convention centre in Singapore (Section 6.4, p. 149). We focussed the evaluation particularly on investigating the system performance and accuracy of the tracking mechanism – conducting over 150 controlled walk-by experiments across five different display locations and three different sizes of spatial fences for detecting the viewer's proximity to a display.

4. We integrated the new lottery scheduling approach part of C3.1 as the new main content scheduling mechanism in e-Campus displays and trialled the system both in laboratory experiments as well as as part of a complete roll-out (Section 6.6, p. 164). In particular, the roll-out consisted of over 65 displays, different lottery ticket allocation approaches and an overview of the system performance and accuracy regarding the scheduling process. We showed that the Lottery Scheduler is technically feasible to serve as a content scheduling system in the context of e-Campus.

7.4 Future Work

Our work has focussed on three major strands of digital signage analytics: the identification and capture of relevant datasets, the use of such data for the creation of novel analytics reports, and supporting the utilisation of analytics for actuation purposes on the screens themselves – ultimately feeding information back into the display network. In this section, we provide an overview of future directions and research from each of the three strands.

7.4.1 The Physical Cookie

The ultimate vision for future digital signage analytics is to support the physical equivalent of a click-through event – being able to tell whether individuals from an audience have followed up on content that has been seen on a display by, for example, purchasing an advertised product or changing their behaviour in other ways. In order to create such detailed analytics, however, accurate and comprehensive tracking of individuals may become necessary, leading to a set of technologies that potentially violate the privacy of individuals. In the Web, the tracking of viewers is often achieved through the use of browser cookies that hold small pieces of information stored by a Website on the user's device and included in a subsequent request to the same Web site. For the Website provider, the advantages of using cookies are the ability to recognise returning visitors of a website and track the visitors' activities and behaviour across subsequent visits across multiple days or even months [MDN18]. For visitors and users of Websites, cookies have the advantage that they can be easily deleted with the consequence that a subsequent visit to the same server makes the user look like a new user. In the physical world, particularly with the use of infrastructure-based and biometric tracking technologies (e.g. facial recognition), an equivalent to 'deleting the cookie' does not currently exist and it becomes impossible for viewers to understand the level of sensing and tracking tacking place in a space. We believe that an opportunity for future work in this space is to consider the design and development of a 'physical cookie' – an equivalent of the browser cookie applied in the 'physical' world providing the viewer more control about the sensing and tracking that takes place regarding the individual. With the use of a cookie model, the tracking and sensing in an environment may become more transparent to the individual and, more crucially, provide a level of control to the extent to which environments track the individual, e.g. when returning to the same space. Whether such a 'physical cookie' is conceptionally modelled as a tangible device, a virtual data container located on the user's device, or stored and maintained 'in the edge' (e.g. in the form of "Privacy Mediators" and Cloudlets [Dav+16]) may be part of the future research in this space.

7.4.2 Synthetic Analytics

We originally introduced the concept of 'synthetic analytics' as means of capturing comprehensive viewer mobility data (Section 3.4.3, p. 76) and creating novel viewer-centric analytics reports. In particular, we presented the approach as a substitute for viewer- and infrastructure-based tracking that preserves viewer privacy and does not reveal sensible insights about individuals, primarily focussing on using the approach as a mechanism to enable us to explore the creation of novel viewer-centric analytics reports without the need to collect sensitive datasets about individuals. We believe that the next step in the exploitation of the synthetic analytics approach may consist of the design and development of mobility models of higher accuracy to strengthen the overall approach. Such mobility models with higher accuracy are particularly crucial for the creation of precise analytics reports – making the synthetic analytics approach feasible to be used as a substitute for the collection of sensible viewer-related data. The creation of mobility models could, for example, be based on observations or existing mobility traces such as the display sighting logs captured as part of the Tacita display personalisation system.

Whilst we have applied our synthetic analytics approach in the context of digital signs where it was crucial to trace individuals within a spatial space and understand the individuals' interactions with displays deployed in the space, we believe that future work may consist of further developing and generalising the overall approach beyond the domain of digital signs. In particular, synthetic analytics may become relevant and beneficial for a range of application domains that rely on capturing and using mobility traces. This may become particularly interesting for spaces in which it is technically (or ethically) infeasible to trace individuals and create comprehensive mobility patterns. For example, synthetic traces can provide insights into user interactions with 'smart spaces' and such environments that are highly equipped with sensing technology – without collecting sensitive analytics about the individual.

7.4.3 Scheduling for the Individual

In state-of-the-art digital signage deployments, content creators and display owners typically create individual content schedules for specific displays or group of displays separately. Whilst this is suitable for maintaining full control over the content that displays are showing at any given point of time, it may become challenging to actually reach the desired target group for certain content. Future work may consist of the design and development of a system that supports the broadcasting of content for specific target groups through a heterogeneous set of devices and modalities. For example, content may be broadcast using a mix of public displays and personal devices such as mobile phones and tablets, depending on the type of information that is distributed. A subset of the analytics techniques and systems presented in this thesis can serve as a foundation allowing stakeholders to reach the desired target group (e.g. using tracking and analytics techniques). Additionally, future work may consist of gaining an understanding as to when and how to deliver and distribute information to individuals. We see the development of a 'scheduling for the individual' approach as a natural continuation of our research that builds on top of existing sensing and analytics technology and ultimately introduces novel means to actuate both private and public displays. The consideration of additional means for distributing content that go beyond public digital signs (i.e. the use of personal devices) provide an opportunity to further explore cross-device analytics and tracking as drivers to gain a better understanding regarding user interactions across the digital and physical world.

7.5 Closing Remarks

We have seen a rapid increase in the number of digital signs and public displays [Mar17] and market reports predict a significant growth of public displays to a total of around 87 million digital signs by 2021 [Dig17; Mar17]. The wide range of application areas of digital signs including way-finding, advertisement billboards, and arrivals and departures boards at airports and railway stations further indicate the importance of digital signs in the urban environment and beyond. Digital signs and public displays are used by a large number of viewers and form, in some cases, a part of the critical infrastructure. With the growing reliance on digital signs, it becomes important to understand how viewers interact with and benefit from both the displays themselves and the content and information that is shown. In particular, capturing and understanding viewer interactions across a network of displays (e.g. while viewers move across an airport or university campus and subsequently engage with a number of displays) will enable display and content providers to identify areas of improvement and gain a better understanding of the overall user experience with a display network.

In this thesis, we introduced a set of tools and techniques that help capture and understand the use of displays. We believe that the work outlined in the context of this thesis lays the foundation for more general people-centric analytics that go beyond the domain of digital signs and will enable unique analytical insights and understanding into how users interact across the physical and digital world.

Bibliography

[Adi+14] Fadel Adib, Zachary Kabelac, Dina Katabi, and Robert C. Miller. "3D Tracking via Body Radio Reflections". In: *Proceedings of the 11th USENIX Conference on Networked Systems Design and Implementation*. NSDI'14. Seattle, WA: USENIX Association, 2014, pp. 317–329. ISBN: 978-1-931971-09-6. URL: http://dl.acm.org/citation.cfm?id=2616448.2616478.

[Air13] AirTight Networks. *AirTight Social Wi-Fi and Analytics for the Retail Store of the Future*. Solution Brief. AirTight Networks, Inc., 2013.

[Alt+15] Florian Alt, Andreas Bulling, Gino Gravanis, and Daniel Buschek. "GravitySpot: Guiding Users in Front of Public Displays Using On-Screen Visual Cues". In: *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology*. UIST '15. Charlotte, NC, USA: ACM, 2015, pp. 47–56. ISBN: 978-1-4503-3779-3. DOI: 10.1145/2807442.2807490. URL: http://doi.acm.org/10.1145/2807442.2807490.

[Alt+11a] Florian Alt, Thomas Kubitza, Dominik Bial, Firas Zaidan, Markus Ortel, Björn Zurmaar, Tim Lewen, Alireza Sahami Shirazi, and Albrecht Schmidt. "Digifieds: Insights into Deploying Digital Public Notice Areas in the Wild". In: *Proceedings of the 10th International Conference on Mobile and Ubiquitous Multimedia*. MUM '11. Beijing, China: ACM, 2011, pp. 165–174. ISBN: 978-1-4503-1096-3. DOI: 10.1145/2107596.2107618. URL: http://doi.acm.org/10.1145/2107596.2107618.

[Alt+11b] Florian Alt, Nemanja Memarovic, Ivan Elhart, Dominik Bial, Albrecht Schmidt, Marc Langheinrich, Gunnar Harboe, Elaine Huang, and Marcello P. Scipioni. "Designing Shared Public Display Networks – Implications from Today's Paper-Based Notice Areas". In: *Pervasive Computing: 9th International Conference, Pervasive 2011, San Francisco, USA, June 12-15, 2011. Proceedings*. Ed. by Kent Lyons, Jeffrey Hightower, and Elaine M. Huang. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 258–275. ISBN: 978-3-642-21726-5. DOI: 10.1007/978-3-642-21726-5_17. URL: http://dx.doi.org/10.1007/978-3-642-21726-5_17.

[AMS12] Florian Alt, Jorg Müller, and Albrecht Schmidt. "Advertising on Public Display Networks". In: *Computer* 45.5 (May 2012), pp. 50–56. ISSN: 0018-9162. DOI: 10.1109/MC.2012.150. URL: http://dx.doi.org/10.1109/MC.2012.150.

[Alt+13] Florian Alt, Stefan Schneegass, Michael Girgis, and Albrecht Schmidt. "Cognitive Effects of Interactive Public Display Applications". In: *Proceedings of the 2Nd ACM International Symposium on Pervasive Displays*. PerDis '13. Mountain View, California: ACM, 2013, pp. 13–18. ISBN: 978-1-4503-2096-2. DOI: 10.1145/2491568.2491572. URL: http://doi.acm.org/10.1145/2491568.2491572.

[Alt+12] Florian Alt, Stefan Schneegaß, Albrecht Schmidt, Jörg Müller, and Nemanja Memarovic. "How to Evaluate Public Displays". In: *Proceedings of the 2012 International Symposium on Pervasive Displays*. PerDis '12. Porto, Portugal: ACM, 2012, 17:1–17:6. ISBN: 978-1-4503-1414-5. DOI: 10.1145/2307798.2307815. URL: http://doi.acm.org/10.1145/2307798.2307815.

[Ant+06] Gianluca Antonini, Santiago Venegas Martinez, Michel Bierlaire, and Jean Philippe Thiran. "Behavioral Priors for Detection and Tracking of Pedestrians in Video Sequences". In: *International Journal of Computer Vision* 69.2 (2006), pp. 159–180. ISSN: 1573-1405. DOI: 10.1007/s11263-005-4797-0. URL: http://dx.doi.org/10.1007/s11263-005-4797-0.

[AML13] Ajaj Aravamudhan, Archan Misra, and Hoong Chuin Lau. "'Network-theoretic' queuing delay estimation in theme park attractions". In: *2013 IEEE International Conference on Automation Science and Engineering (CASE)*. Aug. 2013, pp. 776–782. DOI: 10.1109/CoASE.2013.6653930.

[AA15] Remzi H. Arpaci-Dusseau and Andrea C. Arpaci-Dusseau. *Operating Systems: Three Easy Pieces*. 0.91. Arpaci-Dusseau Books, May 2015.

[BMG10] Till Ballendat, Nicolai Marquardt, and Saul Greenberg. "Proxemic Interaction: Designing for a Proximity and Orientation-aware Environment". In: *ACM International Conference on Interactive Tabletops and Surfaces*. ITS '10. Saarbrücken, Germany: ACM, 2010, pp. 121–130. ISBN: 978-1-4503-0399-6. DOI: 10.1145/1936652.1936676. URL: http://doi.acm.org/10.1145/1936652.1936676.

[Bec+13] Richard Becker, Ramón Cáceres, Karrie Hanson, Sibren Isaacman, Ji Meng Loh, Margaret Martonosi, James Rowland, Simon Urbanek, Alexander Varshavsky, and Chris Volinsky. "Human Mobility Characterization from Cellular Network Data". In: *Commun. ACM* 56.1 (Jan. 2013), pp. 74–82. ISSN: 0001-0782. DOI: 10.1145/2398356.2398375. URL: http://doi.acm.org/10.1145/2398356.2398375.

[BS00] Bettina Berendt and Myra Spiliopoulou. "Analysis of Navigation Behaviour in Web Sites Integrating Multiple Information Systems". In: *The VLDB Journal* 9.1 (Mar. 2000), pp. 56–75. ISSN: 1066-8888. DOI: 10.1007/s007780050083. URL: http://dx.doi.org/10.1007/s007780050083.

[Bon02] Eric Bonabeau. "Agent-based modeling: Methods and techniques for simulating human systems". In: *Proceedings of the National Academy of Sciences* 99.3 (2002), pp. 7280–7287. ISSN: 0027-8424. DOI: 10.1073/pnas.082080899. URL: http://www.pnas.org/content/99/suppl_3/7280.

[BH13] Agata Brajdic and Robert Harle. "Walk Detection and Step Counting on Unconstrained Smartphones". In: *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. UbiComp '13. Zurich, Switzerland: ACM, 2013, pp. 225–234. ISBN: 978-1-4503-1770-2. DOI: 10.1145/2493432.2493449. URL: http://doi.acm.org/10.1145/2493432.2493449.

[Bro] BroadSign Digital Signage Software Solutions. http://broadsign.com/ [Last accessed: February 2015].

[BBW07] Jason Burbie, Angie Brown, and WAA Standards Committee. *Web Analytics Definitions – Version 4.0*. Tech. rep. Web Analytics Association, 2007.

[Bur+05] Moira Burke, Anthony Hornof, Erik Nilsen, and Nicholas Gorman. "High-cost Banner Blindness: Ads Increase Perceived Workload, Hinder Visual Search, and Are Forgotten". In: *ACM Transactions on Computer-Human Interaction* 12.4 (Dec. 2005), pp. 423–445. ISSN: 1073-0516. DOI: 10.1145/1121112.1121116. URL: http://doi.acm.org/10.1145/1121112.1121116.

[Can+10] J. Candamo, M. Shreve, D. B. Goldgof, D. B. Sapper, and R. Kasturi. "Understanding Transit Scenes: A Survey on Human Behavior-Recognition Algorithms". In: *IEEE Transactions on Intelligent Transportation Systems* 11.1 (Mar. 2010), pp. 206–224. ISSN: 1524-9050. DOI: 10.1109/TITS.2009.2030963.

[Cap17] Captive Media Ltd. *Captive Media Advertising*. July 2017. URL: http://www.captive-media.co.uk.

[Car+04] Scott Carter, Elizabeth Churchill, Laurent Denoue, Jonathan Helfman, and Les Nelson. “Digital Graffiti: Public Annotation of Multimedia Content”. In: *CHI ’04 Extended Abstracts on Human Factors in Computing Systems*. CHI EA ’04. Vienna, Austria: ACM, 2004, pp. 1207–1210. ISBN: 1-58113-703-6. DOI: 10.1145/985921.986025. URL: http://doi.acm.org/10.1145/985921.986025.

[Cas+90] J.D. Case, M. Fedor, M.L. Schoffstall, and J. Davin. *Simple Network Management Protocol (SNMP)*. Tech. rep. May 1990. DOI: 10.17487/rfc1157. URL: https://doi.org/10.17487%2Frfc1157.

[Cav11] Ann Cavoukian. *Anonymous Video Analytics (AVA) technology and privacy*. White Paper. Information and Privacy Commissioner, Ontario, Canada, Apr. 2011. URL: http://www.ontla.on.ca/library/repository/mon/25005/309565.pdf.

[CAY] CAYIN. *SuperReporter 2 – Generate Statistics Reports for Billing and Performance Analysis*. Tech. rep. CAYIN.

[Che+13] Shih-Fen Cheng, Larry Lin, Jiali Du, Hoong Chuin Lau, and Pradeep Varakantham. “An agent-based simulation approach to experience management in theme parks”. In: *2013 Winter Simulations Conference (WSC)*. Dec. 2013, pp. 1527–1538. DOI: 10.1109/WSC.2013.6721536. URL: https://doi.org/10.1109/WSC.2013.6721536.

[Che+05] Keith Cheverst, Alan Dix, Daniel Fitton, Chris Kray, Mark Rouncefield, Corina Sas, George Saslis-Lagoudakis, and Jennifer G. Sheridan. “Exploring Bluetooth Based Mobile Phone Interaction with the Hermes Photo Display”. In: *Proceedings of the 7th International Conference on Human Computer Interaction with Mobile Devices &Amp; Services*. MobileHCI ’05. Salzburg, Austria: ACM, 2005, pp. 47–54. ISBN: 1-59593-089-2. DOI: 10.1145/1085777.1085786. URL: http://doi.acm.org/10.1145/1085777.1085786.

[Chu+03] Elizabeth F. Churchill, Les Nelson, Laurent Denoue, and Andreas Girgensohn. “The Plasma Poster Network: Posting Multimedia Content in Public Places”. In: *Proceedings of INTERACT ’03*. Zurich, Switzerland: IOS Press, 2003. URL: http://www.fxpal.com/publications/FXPAL-PR-03-197.pdf.

[CL15] G. D. Clark and J. Lindqvist. “Engineering Gesture-Based Authentication Systems”. In: *IEEE Pervasive Computing* 14.1 (Jan. 2015), pp. 18–25. ISSN: 1536-1268. DOI: 10.1109/MPRV.2015.6.

[Cli13] Sarah Clinch. “Supporting user appropriation of public displays”. English. PhD thesis. Lancaster University, 2013.

[Cli+13] Sarah Clinch, Nigel Davies, Adrian Friday, and Graham Clinch. “Yarely: A Software Player for Open Pervasive Display Networks”. In: *Proceedings of the 2Nd ACM International Symposium on Pervasive Displays*. PerDis ’13. Mountain View, California: ACM, 2013, pp. 25–30. ISBN: 978-1-4503-2096-2. DOI: 10.1145/2491568.2491575. URL: http://doi.acm.org/10.1145/2491568.2491575.

[Cli+14] Sarah Clinch, Mateusz Mikusz, Miriam Greis, Nigel Davies, and Adrian Friday. “Mercury: An Application Store for Open Display Networks”. In: *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. UbiComp ’14. Seattle, Washington: ACM, 2014, pp. 511–522. ISBN: 978-1-4503-2968-2. DOI: 10.1145/2632048.2636080. URL: http://doi.acm.org/10.1145/2632048.2636080.

[CMS99] Robert Cooley, Bamshad Mobasher, and Jaideep Srivastava. “Data Preparation for Mining World Wide Web Browsing Patterns”. In: *Knowledge and Information Systems* 1.1 (Feb. 1999), pp. 5–32. ISSN: 0219-3116. DOI: 10.1007/BF03325089. URL: https://doi.org/10.1007/BF03325089.

[CMS97] Robert Cooley, Bamshad Mobasher, and Jaideep Srivastava. "Web mining: information and pattern discovery on the World Wide Web". In: *Proceedings Ninth IEEE International Conference on Tools with Artificial Intelligence*. Nov. 1997, pp. 558–567. DOI: 10.1109/TAI.1997.632303. URL: https://doi.org/10.1109/TAI.1997.632303.

[Cor] Convergent Media Systems Corporation. *SPINETIX Cockpit*. Tech. rep. Convergent Media Systems Corporation.

[DCM15] Nicholas S. Dalton, Emily Collins, and Paul Marshall. "Display Blindness?: Looking Again at the Visibility of Situated Displays Using Eye-tracking". In: *Proceedings of CHI '15*. Seoul, Republic of Korea: ACM, 2015, pp. 3889–3898. ISBN: 978-1-4503-3145-6. DOI: 10.1145/2702123.2702150.

[DCA14] Nigel Davies, Sarah Clinch, and Florian Alt. "Pervasive Displays: Understanding the Future of Digital Signage". In: *Synthesis Lectures on Mobile and Pervasive Computing* 8.1 (2014), pp. 1–128. DOI: 10.2200/S00558ED1V01Y201312MPC011. URL: https://doi.org/10.2200/S00558ED1V01Y201312MPC011.

[Dav+09] Nigel Davies, Adrian Friday, Peter Newman, Sarah Rutlidge, and Oliver Storz. "Using Bluetooth Device Names to Support Interaction in Smart Environments". In: *Proceedings of the 7th International Conference on Mobile Systems, Applications, and Services*. MobiSys '09. Krakov, Poland: ACM, 2009, pp. 151–164. ISBN: 978-1-60558-566-6. DOI: 10.1145/1555816.1555832. URL: http://doi.acm.org/10.1145/1555816.1555832.

[Dav+14] Nigel Davies, Marc Langheinrich, Sarah Clinch, Ivan Elhart, Adrian Friday, Thomas Kubitza, and Bholanathsingh Surajbali. "Personalisation and Privacy in Future Pervasive Display Networks". In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '14. Toronto, Ontario, Canada: ACM, 2014, pp. 2357–2366. ISBN: 978-1-4503-2473-1. DOI: 10.1145/2556288.2557287. URL: http://doi.acm.org/10.1145/2556288.2557287.

[Dav+12] Nigel Davies, Marc Langheinrich, Rui Jose, and Albrecht Schmidt. "Open Display Networks: A Communications Medium for the 21st Century". In: *Computer* 45.5 (May 2012), pp. 58–64. ISSN: 0018-9162. DOI: 10.1109/MC.2012.114. URL: https://doi.org/10.1109/MC.2012.114.

[Dav+16] Nigel Davies, Nina Taft, Mahadev Satyanarayanan, Sarah Clinch, and Brandon Amos. "Privacy Mediators: Helping IoT Cross the Chasm". In: *Proceedings of the 17th International Workshop on Mobile Computing Systems and Applications*. HotMobile '16. St. Augustine, Florida, USA: ACM, 2016, pp. 39–44. ISBN: 978-1-4503-4145-5. DOI: 10.1145/2873587.2873600. URL: http://doi.acm.org/10.1145/2873587.2873600.

[DWD07] Alexander De Luca, Roman Weiss, and Heiko Drewes. "Evaluation of Eye-gaze Interaction Methods for Security Enhanced PIN-entry". In: *Proceedings of the 19th Australasian Conference on Computer-Human Interaction: Entertaining User Interfaces*. OZCHI '07. Adelaide, Australia: ACM, 2007, pp. 199–202. ISBN: 978-1-59593-872-5. DOI: 10.1145/1324892.1324932. URL: http://doi.acm.org/10.1145/1324892.1324932.

[Dig17] DigitalSignageToday. *Global shipments of digital signage displays to reach 22.9 million by 2021*. 2017. URL: https://www.digitalsignagetoday.com/news/global-shipments-of-digital-signage-displays-to-reach-229-million-by-2021/.

[Dob+11] Florin Dobrian, Vyas Sekar, Asad Awan, Ion Stoica, Dilip Joseph, Aditya Ganjam, Jibin Zhan, and Hui Zhang. "Understanding the Impact of Video Quality on User Engagement". In: *Proceedings of the ACM SIGCOMM 2011 Conference*. SIGCOMM '11. Toronto, Ontario, Canada: ACM, 2011, pp. 362–373. ISBN: 978-1-4503-0797-0. DOI: 10.1145/2018436.2018478. URL: http://doi.acm.org/10.1145/2018436.2018478.

[Dos+14] Jakub Dostal, Uta Hinrichs, Per Ola Kristensson, and Aaron Quigley. "SpiderEyes: Designing Attention- and Proximity-aware Collaborative Interfaces for Wall-sized Displays". In: *Proceedings of the 19th International Conference on Intelligent User Interfaces*. IUI '14. Haifa, Israel: ACM, 2014, pp. 143–152. ISBN: 978-1-4503-2184-6. DOI: 10.1145/2557500.2557541. URL: http://doi.acm.org/10.1145/2557500.2557541.

[Ead+04] M. Eaddy, G. Blasko, J. Babcock, and S. Feiner. "My own private kiosk: privacy-preserving public displays". In: *Eighth International Symposium on Wearable Computers*. Vol. 1. Oct. 2004, pp. 132–135. DOI: 10.1109/ISWC.2004.32.

[Eds14] Jr Edson C Tandoc. "Journalism is twerking? How web analytics is changing the process of gatekeeping". In: *New Media & Society* 16.4 (2014), pp. 559–575. DOI: 10.1177/1461444814530541. eprint: https://doi.org/10.1177/1461444814530541. URL: https://doi.org/10.1177/1461444814530541.

[Elh+14] Ivan Elhart, Marc Langheinrich, Nemanja Memarovic, and Tommi Heikkinen. "Scheduling Interactive and Concurrently Running Applications in Pervasive Display Networks". In: *Proceedings of The International Symposium on Pervasive Displays*. PerDis '14. Copenhagen, Denmark: ACM, 2014, 104:104–104:109. ISBN: 978-1-4503-2952-1. DOI: 10.1145/2611009.2611039. URL: http://doi.acm.org/10.1145/2611009.2611039.

[Elh+17] Ivan Elhart, Mateusz Mikusz, Cristian Gomez Mora, Marc Langheinrich, and Nigel Davies. "Audience Monitor: An Open Source Tool for Tracking Audience Mobility in Front of Pervasive Displays". In: *Proceedings of the 6th ACM International Symposium on Pervasive Displays*. PerDis '17. Lugano, Switzerland: ACM, 2017, 10:1–10:8. ISBN: 978-1-4503-5045-7. DOI: 10.1145/3078810.3078823. URL: http://doi.acm.org/10.1145/3078810.3078823.

[Esp17] Esprida Corporation. *Esprida Smart Monitoring*. 2017. URL: http://espridalivecontrol.com/tour/smart-monitoring (visited on 04/25/2017).

[Fan+09] Quanfu Fan, A. Yanagawa, R. Bobbitt, Yun Zhai, R. Kjeldsen, S. Pankanti, and A. Hampapur. "Fast detection of retail fraud using polar touch buttons". In: *Multimedia and Expo, 2009. ICME 2009. IEEE International Conference on*. June 2009, pp. 1266–1269. DOI: 10.1109/ICME.2009.5202732. URL: http://dx.doi.org/10.1109/ICME.2009.5202732.

[Far+14] Giovanni Maria Farinella, Giuseppe Farioli, Sebastiano Battiato, Salvo Leonardi, and Giovanni Gallo. "Face Re-Identification for Digital Signage Applications". In: *Video Analytics for Audience Measurement: First International Workshop, VAAM 2014, Stockholm, Sweden, August 24, 2014. Revised Selected Papers*. Ed. by Cosimo Distante, Sebastiano Battiato, and Andrea Cavallaro. Cham: Springer International Publishing, 2014, pp. 40–52. ISBN: 978-3-319-12811-5. DOI: 10.1007/978-3-319-12811-5_3. URL: http://dx.doi.org/10.1007/978-3-319-12811-5_3.

[Fin+96] Joe Finney, Stephen Wade, Nigel Davies, and Adrian Friday. "Flump: The FLexible Ubiquitous Monitor Project". In: *Cabernet Radicals Workshop*. May 1996. URL: http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.26.2692.

[Fraa] Fraunhofer IIS. *Fraunhofer AVARD*. http://www.iis.fraunhofer.de/en/ff/bsy/tech/bildanalyse/avard.html. (Visited on 10/22/2018).

[Frab] Fraunhofer IIS. *Fraunhofer SHORE*. www.iis.fraunhofer.de/en/ff/bsy/tech/bildanalyse/shore-gesichtsdetektion.html. (Visited on 10/22/2018).

[FDE12] Adrian Friday, Nigel Davies, and Christos Efstratiou. "Reflections on Long-Term Experiments with Public Displays". In: *Computer* 45.5 (May 2012), pp. 34–41. ISSN: 0018-9162. DOI: 10.1109/MC.2012.155. URL: https://doi.org/10.1109/MC.2012.155.

[GH15] Amir Gandomi and Murtaza Haider. “Beyond the hype: Big data concepts, methods, and analytics”. In: *International Journal of Information Management* 35.2 (2015), pp. 137–144. ISSN: 0268-4012. DOI: http://dx.doi.org/10.1016/j.ijinfomgt.2014.10.007. URL: http://www.sciencedirect.com/science/article/pii/S0268401214001066.

[GHN13] Matthias Geel, Daniel Huguenin, and Moira C. Norrie. “PresiShare: Opportunistic Sharing and Presentation of Content Using Public Displays and QR Codes”. In: *Proceedings of the 2Nd ACM International Symposium on Pervasive Displays*. PerDis ’13. Mountain View, California: ACM, 2013, pp. 103–108. ISBN: 978-1-4503-2096-2. DOI: 10.1145/2491568.2491591. URL: http://doi.acm.org/10.1145/2491568.2491591.

[GF08] Anupam Ghosh and Jane Fedorowicz. “The role of trust in supply chain governance”. In: *Business Process Management Journal* 14.4 (2008), pp. 453–470. DOI: 10.1108/14637150810888019. URL: http://dx.doi.org/10.1108/14637150810888019.

[Gil+14] Nicholas Gillian, Sara Pfenninger, Spencer Russell, and Joseph A. Paradiso. “Gestures Everywhere: A Multimodal Sensor Fusion and Analysis Framework for Pervasive Displays”. In: *Proceedings of The International Symposium on Pervasive Displays*. PerDis ’14. Copenhagen, Denmark: ACM, 2014, 98:98–98:103. ISBN: 978-1-4503-2952-1. DOI: 10.1145/2611009.2611032. URL: http://doi.acm.org/10.1145/2611009.2611032.

[GSW08] Andreas Girgensohn, Frank Shipman, and Lynn Wilcox. “Determining Activity Patterns in Retail Spaces Through Video Analysis”. In: *Proceedings of the 16th ACM International Conference on Multimedia*. MM ’08. Vancouver, British Columbia, Canada: ACM, 2008, pp. 889–892. ISBN: 978-1-60558-303-7. DOI: 10.1145/1459359.1459514. URL: http://doi.acm.org/10.1145/1459359.1459514.

[Glo17] Global Market Insights. *Digital Signage Market size to exceed $23bn by 2023*. 2017. URL: https://www.digitalsignagetoday.com/news/global-shipments-of-digital-signage-displays-to-reach-229-million-by-2021/ (visited on 10/22/2018).

[GR08] Bruno Gonçalves and José Ramasco. “Human dynamics revealed through Web analytics”. In: *Phys. Rev. E* 78 (2 Aug. 2008), p. 026123. DOI: 10.1103/PhysRevE.78.026123. URL: https://link.aps.org/doi/10.1103/PhysRevE.78.026123.

[Goo18a] Google Inc. *About goals – Analytics Help*. 2018. URL: https://support.google.com/analytics/answer/1012040?hl=en (visited on 06/28/2018).

[Goo18b] Google Inc. *Bounce rate – Analytics Help*. 2018. URL: https://support.google.com/analytics/answer/1009409?hl=en (visited on 06/28/2018).

[Goo17] Google Inc. *DoubleClick by Google*. 2017. URL: https://www.doubleclick.com (visited on 09/01/2017).

[Goo18c] Google Inc. *Google Analytics – About Events*. 2018. URL: https://support.google.com/analytics/answer/1033068 (visited on 09/01/2017).

[Goo18d] Google Inc. *Google Analytics – Set up Analytics tracking with gtag.js*. 2018. URL: https://support.google.com/analytics/answer/2763052?hl=en (visited on 03/22/2019).

[Goo18e] Google Inc. *Google Analytics – Set up Analytics tracking with gtag.js*. 2018. URL: https://support.google.com/analytics/answer/1008080 (visited on 02/08/2018).

[Goo18f] Google Inc. *Google Analytics – Using the flow visualization reports*. 2018. URL: https://support.google.com/analytics/answer/2519989 (visited on 09/01/2017).

[Goo18g] Google Inc. *How a web session is defined in Analytics*. 2018. URL: https://support.google.com/analytics/answer/2731565 (visited on 06/28/2018).

[Goo18h] Google Inc. *Measurement Protocol Parameter Reference*. https://developers.google.com/analytics/devguides/collection/protocol/v1/. Feb. 2018.

[Goo18i] Google Inc. *Pageview – Analytics Help*. 2018. URL: https://support.google.com/analytics/answer/6086080?hl=en (visited on 06/28/2018).

[Gra+13] Kazjon Grace, Rainer Wasinger, Christopher Ackad, Anthony Collins, Oliver Dawson, Richard Gluga, Judy Kay, and Martin Tomitsch. "Conveying Interactivity at an Interactive Public Information Display". In: *Proceedings of the 2Nd ACM International Symposium on Pervasive Displays*. PerDis '13. Mountain View, California: ACM, 2013, pp. 19–24. ISBN: 978-1-4503-2096-2. DOI: 10.1145/2491568.2491573. URL: http://doi.acm.org/10.1145/2491568.2491573.

[GJC04] M.K. Gray, J.P.I.I. Jeffrey, and Y. Chery. *Position detection and location tracking in a wireless network*. Jan. 2004. URL: https://www.google.com/patents/US6674403.

[Gre+14] Saul Greenberg, Sebastian Boring, Jo Vermeulen, and Jakub Dostal. "Dark Patterns in Proxemic Interactions: A Critical Perspective". In: *Proceedings of the 2014 Conference on Designing Interactive Systems*. DIS '14. Vancouver, BC, Canada: ACM, 2014, pp. 523–532. ISBN: 978-1-4503-2902-6. DOI: 10.1145/2598510.2598541. URL: http://doi.acm.org/10.1145/2598510.2598541.

[Gre+11] Saul Greenberg, Nicolai Marquardt, Till Ballendat, Rob Diaz-Marino, and Miaosen Wang. "Proxemic Interactions: The New Ubicomp?" In: *interactions* 18.1 (Jan. 2011), pp. 42–50. ISSN: 1072-5520. DOI: 10.1145/1897239.1897250. URL: http://doi.acm.org/10.1145/1897239.1897250.

[Ham+09] Arun Hampapur, Russell Bobbitt, Lisa Brown, Mike Desimone, Rogerio Feris, Rick Kjeldsen, Max Lu, Carl Mercier, Chris Milite, Stephen Russo, Chiao-fe Shu, and Yun Zhai. "Video Analytics in Urban Environments". In: *2009 Sixth IEEE International Conference on Advanced Video and Signal Based Surveillance*. Sept. 2009, pp. 128–133. DOI: 10.1109/AVSS.2009.100. URL: https://doi.org/10.1109/AVSS.2009.100.

[Han+13] J. Han, L. Shao, D. Xu, and J. Shotton. "Enhanced Computer Vision With Microsoft Kinect Sensor: A Review". In: *IEEE Transactions on Cybernetics* 43.5 (Oct. 2013), pp. 1318–1334. ISSN: 2168-2267. DOI: 10.1109/TCYB.2013.2265378. URL: https://doi.org/10.1109/TCYB.2013.2265378.

[HRD11] John Hardy, Enrico Rukzio, and Nigel Davies. "Real World Responses to Interactive Gesture Based Public Displays". In: *Proceedings of the 10th International Conference on Mobile and Ubiquitous Multimedia*. MUM '11. Beijing, China: ACM, 2011, pp. 33–39. ISBN: 978-1-4503-1096-3. DOI: 10.1145/2107596.2107600. URL: http://doi.acm.org/10.1145/2107596.2107600.

[Hu+13] Pan Hu, Liqun Li, Chunyi Peng, Guobin Shen, and Feng Zhao. "Pharos: Enable Physical Analytics Through Visible Light Based Indoor Localization". In: *Proceedings of the Twelfth ACM Workshop on Hot Topics in Networks*. HotNets-XII. College Park, Maryland: ACM, 2013, 5:1–5:7. ISBN: 978-1-4503-2596-7. DOI: 10.1145/2535771.2535790. URL: http://doi.acm.org/10.1145/2535771.2535790.

[HKB08] Elaine M. Huang, Anna Koster, and Jan Borchers. "Overcoming Assumptions and Uncovering Practices: When Does the Public Really Look at Public Displays?" In: *Proceedings of Pervasive '08*. Sydney, Australia: Springer, 2008. ISBN: 978-3-540-79575-9. DOI: 10.1007/978-3-540-79576-6_14. URL: http://dx.doi.org/10.1007/978-3-540-79576-6_14.

[IBM13] IBM Corporation. *IBM Intelligent Video Analytics*. Tech. rep. IBM Corporation, June 2013.

[Int11] Intel Corporation. *A Report on a Field Trial of Anonymous Video Analytics (AVA) in Digital Signage*. Tech. rep. Intel Corporation, 2011.

[Int18] Intel Corporation. *Intel AIM Suite*. https://aimsuite.intel.com/. 2018. (Visited on 10/22/2018).

[Jac+10] Giulio Jacucci, Ann Morrison, Gabriela T. Richard, Jari Kleimola, Peter Peltonen, Lorenza Parisi, and Toni Laitinen. "Worlds of Information: Designing for Engagement at a Public Multi-touch Display". In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '10. Atlanta, Georgia, USA: ACM, 2010, pp. 2267–2276. ISBN: 978-1-60558-929-9. DOI: 10.1145/1753326.1753669. URL: http://doi.acm.org/10.1145/1753326.1753669.

[Jan09] Bernard J. (Jim) Jansen. "Understanding User-Web Interactions via Web Analytics". In: *Synthesis Lectures on Information Concepts, Retrieval, and Services* 1.1 (Jan. 2009), pp. 1–102. DOI: 10.2200/s00191ed1v01y200904icr006. URL: https://doi.org/10.2200/s00191ed1v01y200904icr006.

[Jay+16] Kasthuri Jayarajah, Rajesh Krishna Balan, Meera Radhakrishnan, Archan Misra, and Youngki Lee. "LiveLabs: Building In-Situ Mobile Sensing & Behavioural Experimentation TestBeds". In: *Proceedings of the 14th Annual International Conference on Mobile Systems, Applications, and Services*. MobiSys '16. Singapore, Singapore: ACM, 2016, pp. 1–15. ISBN: 978-1-4503-4269-8. DOI: 10.1145/2906388.2906400. URL: http://doi.acm.org/10.1145/2906388.2906400.

[Jur+13] Marko Jurmu, Masaki Ogawa, Sebastian Boring, Jukka Riekki, and Hideyuki Tokuda. "Waving to a Touch Interface: Descriptive Field Study of a Multipurpose Multimodal Public Display". In: *Proceedings of the 2Nd ACM International Symposium on Pervasive Displays*. PerDis '13. Mountain View, California: ACM, 2013, pp. 7–12. ISBN: 978-1-4503-2096-2. DOI: 10.1145/2491568.2491571. URL: http://doi.acm.org/10.1145/2491568.2491571.

[Kha+13] Azeem J. Khan, Vikash Ranjan, Trung-Tuan Luong, Rajesh Balan, and Archan Misra. "Experiences with performance tradeoffs in practical, continuous indoor localization". In: *2013 IEEE 14th International Symposium on "A World of Wireless, Mobile and Multimedia Networks" (WoWMoM)*. June 2013, pp. 1–9. DOI: 10.1109/WoWMoM.2013.6583387. URL: https://doi.org/10.1109/WoWMoM.2013.6583387.

[Kim+10] David Kim, Paul Dunphy, Pam Briggs, Jonathan Hook, John W. Nicholson, James Nicholson, and Patrick Olivier. "Multi-touch Authentication on Tabletops". In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '10. Atlanta, Georgia, USA: ACM, 2010, pp. 1093–1102. ISBN: 978-1-60558-929-9. DOI: 10.1145/1753326.1753489. URL: http://doi.acm.org/10.1145/1753326.1753489.

[Kon+13] Divyan Munirathnam Konidala, Robert H. Deng, Yingjiu Li, Hoong Chuin Lau, and Stephen E. Fienberg. "Anonymous Authentication of Visitors for Mobile Crowd Sensing at Amusement Parks". In: *Information Security Practice and Experience: 9th International Conference, ISPEC 2013, Lanzhou, China, May 12-14, 2013. Proceedings*. Ed. by Robert H. Deng and Tao Feng. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 174–188. ISBN: 978-3-642-38033-4. DOI: 10.1007/978-3-642-38033-4_13. URL: http://dx.doi.org/10.1007/978-3-642-38033-4_13.

[KB00] Raymond Kosala and Hendrik Blockeel. "Web Mining Research: A Survey". In: *SIGKDD Explor. Newsl.* 2.1 (June 2000), pp. 1–15. ISSN: 1931-0145. DOI: 10.1145/360402.360406. URL: http://doi.acm.org/10.1145/360402.360406.

[Kra+05] N. Krahnstoever, J. Rittscher, P. Tu, K. Chean, and T. Tomlinson. “Activity Recognition using Visual Tracking and RFID”. In: *Application of Computer Vision, 2005. WACV/MOTIONS ’05 Volume 1. Seventh IEEE Workshops on*. Vol. 1. Jan. 2005, pp. 494–500. DOI: 10.1109/ACVMOT.2005.17.

[KGR08] Christian Kray, Areti Galani, and Michael Rohs. “Facilitating Opportunistic Interaction with Ambient Displays”. In: *Workshop on Designing and Evaluating Mobile Phone-Based Interaction with Public Displays at CHI 2008*. Florence, Italy, 2008.

[KW09] Balachander Krishnamurthy and Craig Wills. “Privacy Diffusion on the Web: A Longitudinal Perspective”. In: *Proceedings of the 18th International Conference on World Wide Web*. WWW ’09. Madrid, Spain: ACM, 2009, pp. 541–550. ISBN: 978-1-60558-487-4. DOI: 10.1145/1526709.1526782. URL: http://doi.acm.org/10.1145/1526709.1526782.

[KSO11] Antonio Krüger, Johannes Schöning, and Patrick Olivier. “How Computing Will Change the Face of Retail”. In: *Computer* 44.4 (Apr. 2011), pp. 84–87. ISSN: 0018-9162. DOI: 10.1109/MC.2011.112.

[KLO04] Susan Cohen Kulp, Hau L. Lee, and Elie Ofek. “Manufacturer Benefits from Information Integration with Retail Customers”. In: *Management Science* 50.4 (2004), pp. 431–444. DOI: 10.1287/mnsc.1030.0182. URL: http://dx.doi.org/10.1287/mnsc.1030.0182.

[Lan02] Marc Langheinrich. “A Privacy Awareness System for Ubiquitous Computing Environments”. In: *UbiComp 2002: Ubiquitous Computing*. Ed. by Gaetano Borriello and Lars Erik Holmquist. Berlin, Heidelberg: Springer Berlin Heidelberg, 2002, pp. 237–245. ISBN: 978-3-540-45809-8.

[LS18] Marc Langheinrich and Florian Schaub. “Privacy in Mobile and Pervasive Computing”. In: *Synthesis Lectures on Mobile and Pervasive Computing* 10.1 (Dec. 2018), pp. 1–139. DOI: 10.2200/s00882ed1v01y201810mpc013. URL: https://doi.org/10.2200/s00882ed1v01y201810mpc013.

[LPW04] Hau L. Lee, V. Padmanabhan, and Seungjin Whang. “Information Distortion in a Supply Chain: The Bullwhip Effect”. In: *Management Science* 50.12_supplement (2004), pp. 1875–1886. DOI: 10.1287/mnsc.1040.0266. URL: http://dx.doi.org/10.1287/mnsc.1040.0266.

[LW00] Hau L. Lee and Seungjin Whang. “Information sharing in a supply chain”. In: *International Journal of Manufacturing Technology and Management* 1.1 (2000), pp. 79–93. DOI: 10.1504/IJMTM.2000.001329. URL: http://www.inderscienceonline.com/doi/abs/10.1504/IJMTM.2000.001329.

[LSC12] S. Lee, H. Shin, and H. Cha. “A pedestrian tracking system using group mobility information”. In: *MILCOM 2012 - 2012 IEEE Military Communications Conference*. Oct. 2012, pp. 1–6. DOI: 10.1109/MILCOM.2012.6415710. URL: https://doi.org/10.1109/MILCOM.2012.6415710.

[LB14] Youngki Lee and Rajesh Krishina Balan. “The Case for Human-centric Personal Analytics”. In: *Proceedings of the 2014 Workshop on Physical Analytics*. WPA ’14. Bretton Woods, New Hampshire, USA: ACM, 2014, pp. 25–29. ISBN: 978-1-4503-2825-8. DOI: 10.1145/2611264.2611267. URL: http://doi.acm.org/10.1145/2611264.2611267.

[Leh+12] Janette Lehmann, Mounia Lalmas, Elad Yom-Tov, and Georges Dupret. “Models of User Engagement”. In: *User Modeling, Adaptation, and Personalization: 20th International Conference, UMAP 2012, Montreal, Canada, July 16-20, 2012. Proceedings*. Ed. by Judith Masthoff, Bamshad Mobasher, Michel C. Desmarais, and Roger Nkambou. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 164–175. ISBN: 978-3-642-31454-4. DOI: 10.1007/978-3-642-31454-4_14. URL: http://dx.doi.org/10.1007/978-3-642-31454-4_14.

[LG 13] LG Electronics. *Suntec Singapore's Entrance Gets Lit Up with An Impressive Mega Full HD LED Digital Wall*. 2013. URL: http://www.lg.com/sg/press-release/suntec-singapore-entrance-gets-lit-up-with-an-impressive-mega-full-hd-led-digital-wall (visited on 10/22/2018).

[LCK13] H. Liu, S. Chen, and N. Kubota. "Intelligent Video Systems and Analytics: A Survey". In: *IEEE Transactions on Industrial Informatics* 9.3 (Aug. 2013), pp. 1222–1233. ISSN: 1551-3203. DOI: 10.1109/TII.2013.2255616.

[Loo17] Look Ltd. *Look Digital Signage – Content Manager*. 2017. URL: https://lookdigitalsignage.com/content-manager.html (visited on 04/25/2017).

[Mal+16] Eve Maler, Jean Paoli, C. M. Sperberg-McQueen, François Yergeau, and Tim Bray. *Bluetooth Core Specification v5.0*. Specification of the Bluetooth System. Bluetooth SIG Proprietary, Dec. 2016.

[Mar+11] Nicolai Marquardt, Robert Diaz-Marino, Sebastian Boring, and Saul Greenberg. "The Proximity Toolkit: Prototyping Proxemic Interactions in Ubiquitous Computing Ecologies". In: *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology*. UIST '11. Santa Barbara, California, USA: ACM, 2011, pp. 315–326. ISBN: 978-1-4503-0716-1. DOI: 10.1145/2047196.2047238. URL: http://doi.acm.org/10.1145/2047196.2047238.

[MG12] Nicolai Marquardt and Saul Greenberg. "Informing the Design of Proxemic Interactions". In: *IEEE Pervasive Computing* 11.2 (Feb. 2012), pp. 14–23. ISSN: 1536-1268. DOI: 10.1109/MPRV.2012.15. URL: https://doi.org/10.1109/MPRV.2012.15.

[MG15] Nicolai Marquardt and Saul Greenberg. "Proxemic Interactions: From Theory to Practice". In: *Synthesis Lectures on Human-Centered Informatics* 8.1 (Feb. 2015), pp. 1–199. DOI: 10.2200/s00619ed1v01y201502hci025. URL: https://doi.org/10.2200/s00619ed1v01y201502hci025.

[MZH99] Jean-Philippe Martin-Flatin, Simon Znaty, and Jean-Pierre Hubaux. "A Survey of Distributed Enterprise Network and Systems Management Paradigms". In: *Journal of Network and Systems Management* 7.1 (Mar. 1999), pp. 9–26. ISSN: 1573-7705. DOI: 10.1023/A:1018761615354. URL: https://doi.org/10.1023/A:1018761615354.

[Mar17] Chuck Martin. *38 Million Connected Digital Signs In Use, Heading To 87 Million*. 2017. URL: https://www.mediapost.com/publications/article/297163/38-million-connected-digital-signs-in-use-heading.html (visited on 10/22/2018).

[MR97] Nobuyuki Matsushita and Jun Rekimoto. "HoloWall: Designing a Finger, Hand, Body, and Object Sensitive Wall". In: *Proceedings of the 10th Annual ACM Symposium on User Interface Software and Technology*. UIST '97. Banff, Alberta, Canada: ACM, 1997, pp. 209–210. ISBN: 0-89791-881-9. DOI: 10.1145/263407.263549. URL: http://doi.acm.org/10.1145/263407.263549.

[MCL01] Joseph F. McCarthy, Tony J. Costa, and Edy S. Liongosari. "UniCast, OutCast & GroupCast: Three Steps Toward Ubiquitous, Peripheral Displays". In: *Proceedings of Ubicomp '01*. Atlanta, GA, USA: Springer-Verlag, 2001, pp. 332–345. DOI: 10.1007/3-540-45427-6_28. URL: https://link.springer.com/chapter/10.1007/3-540-45427-6_28.

[MDN18] MDN web docs. *HTTP Cookies*. 2018. URL: https://developer.mozilla.org/en-US/docs/Web/HTTP/Cookies (visited on 01/21/2018).

[MM11] Daniel Michelis and Jörg Müller. "The Audience Funnel: Observations of Gesture Based Interaction With Multiple Large Displays in a City Center". In: *International Journal of Human–Computer Interaction* 27.6 (2011), pp. 562–579. DOI: 10.1080/10447318.2011.555299. URL: http://dx.doi.org/10.1080/10447318.2011.555299.

[Mik+18a] Mateusz Mikusz, Kenny Tsu Wei Choo, Rajesh Krishna Balan, Nigel Davies, and Youngki Lee. "New Challenges in Saturated Displays Environments". In: *IEEE Pervasive Computing* (2018).

[MCD15] Mateusz Mikusz, Sarah Clinch, and Nigel Davies. "Are You Feeling Lucky?: Lottery-based Scheduling for Public Displays". In: *Proceedings of the 4th International Symposium on Pervasive Displays*. PerDis '15. Saarbruecken, Germany: ACM, 2015, pp. 123–129. ISBN: 978-1-4503-3608-6. DOI: 10.1145/2757710.2757721. URL: http://doi.acm.org/10.1145/2757710.2757721.

[MCD17] Mateusz Mikusz, Sarah Clinch, and Nigel Davies. "Design Considerations for Multi-stakeholder Display Analytics". In: *Proceedings of the 6th ACM International Symposium on Pervasive Displays*. PerDis '17. Lugano, Switzerland: ACM, 2017, 18:1–18:10. ISBN: 978-1-4503-5045-7. DOI: 10.1145/3078810.3078830. URL: http://doi.acm.org/10.1145/3078810.3078830.

[Mik+15] Mateusz Mikusz, Sarah Clinch, Rachel Jones, Michael Harding, Christopher Winstanley, and Nigel Davies. "Repurposing Web Analytics to Support the IoT". In: *Computer* 48.9 (Sept. 2015), pp. 42–49. ISSN: 0018-9162. DOI: 10.1109/MC.2015.260. URL: http://doi.org/10.1109/MC.2015.260.

[Mik+18b] Mateusz Mikusz, Sarah Clinch, Peter Shaw, Nigel Davies, and Petteri Nurmi. "Using Pervasive Displays to Aid Student Recall -Reflections on a Campus-Wide Trial". In: *Proceedings of the 7th ACM International Symposium on Pervasive Displays*. PerDis '18. Munich, Germany: ACM, 2018, 6:1–6:8. ISBN: 978-1-4503-5765-4. DOI: 10.1145/3205873.3205882. URL: http://doi.acm.org/10.1145/3205873.3205882.

[Mik+18c] Mateusz Mikusz, Steven Houben, Nigel Davies, Klaus Moessner, and Marc Langheinrich. "Raising awareness of IoT sensor deployments". In: *Living in the Internet of Things: Cybersecurity of the IoT - 2018*. Mar. 2018, pp. 1–8. DOI: 10.1049/cp.2018.0009. URL: https://ieeexplore.ieee.org/document/8379696.

[Mik+16] Mateusz Mikusz, Anastasios Noulas, Nigel Davies, Sarah Clinch, and Adrian Friday. "Next Generation Physical Analytics for Digital Signage". In: *Proceedings of the 3rd International on Workshop on Physical Analytics*. WPA '16. Singapore, Singapore: ACM, 2016, pp. 19–24. ISBN: 978-1-4503-4328-2. DOI: 10.1145/2935651.2935658. URL: http://doi.acm.org/10.1145/2935651.2935658.

[Mik+18d] Mateusz Mikusz, Peter Shaw, Nigel Davies, Sarah Clinch, Ludwig Trotter, Ivan Elhart, Marc Langheinrich, and Adrian Friday. "Experiences of Mobile Personalisation of Pervasive Displays". In: *ACM Transactions on Computer-Human Interaction – TOCHI (in preparation)* (2018).

[MB13] Archan Misra and Rajesh Krishna Balan. "LiveLabs: Initial Reflections on Building a Large-scale Mobile Behavioral Experimentation Testbed". In: *SIGMOBILE Mob. Comput. Commun. Rev.* 17.4 (Dec. 2013), pp. 47–59. ISSN: 1559-1662. DOI: 10.1145/2557968.2557975. URL: http://doi.acm.org/10.1145/2557968.2557975.

[Mit+10] Dennis Mitzel, Esther Horbert, Andreas Ess, and Bastian Leibe. "Multi-person Tracking with Sparse Detection and Continuous Segmentation". In: *Computer Vision – ECCV 2010: 11th European Conference on Computer Vision, Heraklion, Crete, Greece, September 5-11, 2010, Proceedings, Part I*. Ed. by Kostas Daniilidis, Petros Maragos, and Nikos Paragios. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 397–410. ISBN: 978-3-642-15549-9. DOI: 10.1007/978-3-642-15549-9_29. URL: http://dx.doi.org/10.1007/978-3-642-15549-9_29.

[MCS00] Bamshad Mobasher, Robert Cooley, and Jaideep Srivastava. "Automatic Personalization Based on Web Usage Mining". In: *Communications of the ACM* 43.8 (Aug. 2000), pp. 142–151. ISSN: 0001-0782. DOI: 10.1145/345124.345169. URL: http://doi.acm.org/10.1145/345124.345169.

[Mül+12] Jörg Müller, Robert Walter, Gilles Bailly, Michael Nischt, and Florian Alt. "Looking Glass: A Field Study on Noticing Interactivity of a Shop Window". In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '12. Austin, Texas, USA: ACM, 2012, pp. 297–306. ISBN: 978-1-4503-1015-4. DOI: 10.1145/2207676.2207718. URL: http://doi.acm.org/10.1145/2207676.2207718.

[Mül+09] Jörg Müller, Dennis Wilmsmann, Juliane Exeler, Markus Buzeck, Albrecht Schmidt, Tim Jay, and Antonio Krüger. "Display Blindness: The Effect of Expectations on Attention towards Digital Signage". In: *Pervasive Computing: 7th International Conference, Pervasive 2009, Nara, Japan, May 11-14, 2009. Proceedings*. Ed. by Hideyuki Tokuda, Michael Beigl, Adrian Friday, A. J. Bernheim Brush, and Yoshito Tobe. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 1–8. ISBN: 978-3-642-01516-8. DOI: 10.1007/978-3-642-01516-8_1. URL: https://doi.org/10.1007/978-3-642-01516-8_1.

[MFD03] G. Myles, A. Friday, and N. Davies. "Preserving privacy in environments with location-based applications". In: *IEEE Pervasive Computing* 2.1 (Jan. 2003), pp. 56–64. ISSN: 1536-1268. DOI: 10.1109/MPRV.2003.1186726.

[Nan+13] Rajalakshmi Nandakumar, Swati Rallapalli, Krishna Chintalapudi, Venkat Padmanabhan, Lili Qiu, Aishwarya Ganesan, Saikat Guha, Deepanker Aggarwal, and Aakash Goenka. *Physical Analytics: A New Frontier for (Indoor) Location Research*. Tech. rep. Oct. 2013. URL: https://www.microsoft.com/en-us/research/publication/physical-analytics-a-new-frontier-for-indoor-location-research/.

[NEC13] NEC Display Solutions. *NEC FieldAnalyst – Audience Measurement Software*. Product Brochure. NEC Display Solutions Europe GmbH, Apr. 2013.

[Oja+12] T. Ojala, V. Kostakos, H. Kukka, T. Heikkinen, T. Linden, M. Jurmu, S. Hosio, F. Kruger, and D. Zanni. "Multipurpose Interactive Public Displays in the Wild: Three Years Later". In: *Computer* 45.5 (May 2012), pp. 42–49. ISSN: 0018-9162. DOI: 10.1109/MC.2012.115. URL: https://doi.org/10.1109/MC.2012.115.

[OnS17] OnSign TV. *OnSign TV Key Features*. 2017. URL: https://onsign.tv/features/ (visited on 04/25/2017).

[OK14] Morin Ostkamp and Christian Kray. "Supporting Design, Prototyping, and Evaluation of Public Display Systems". In: *Proceedings of the 2014 ACM SIGCHI Symposium on Engineering Interactive Computing Systems*. EICS '14. Rome, Italy: ACM, 2014, pp. 263–272. ISBN: 978-1-4503-2725-1. DOI: 10.1145/2607023.2607035. URL: http://doi.acm.org/10.1145/2607023.2607035.

[Pan+07] Xiaoshan Pan, Charles S. Han, Ken Dauber, and Kincho H. Law. "A multi-agent based framework for the simulation of human and social behaviors during emergency evacuations". In: *AI & SOCIETY* 22.2 (Nov. 2007), pp. 113–132. ISSN: 1435-5655. DOI: 10.1007/s00146-007-0126-1. URL: https://doi.org/10.1007/s00146-007-0126-1.

[PI77] S. S. Panwalkar and Wafik Iskander. "A Survey of Scheduling Rules". In: *Operations Research* 25.1 (1977), pp. 45–61. DOI: 10.1287/opre.25.1.45. eprint: https://doi.org/10.1287/opre.25.1.45. URL: https://doi.org/10.1287/opre.25.1.45.

[PSC13] Kwanghyo Park, Hyojeong Shin, and Hojung Cha. "Smartphone-based pedestrian tracking in indoor corridor environments". In: *Personal and Ubiquitous Computing* 17.2 (2013), pp. 359–370. DOI: 10.1007/s00779-011-0499-5. URL: http://dx.doi.org/10.1007/s00779-011-0499-5.

[PTK18] Callum Parker, Martin Tomitsch, and Judy Kay. "Does the Public Still Look at Public Displays?: A Field Observation of Public Displays in the Wild". In: *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 2.2 (July 2018), 73:1–73:24. ISSN: 2474-9567. DOI: 10.1145/3214276. URL: http://doi.acm.org/10.1145/3214276.

[PKD14] Gonzalo Parra, Joris Klerkx, and Erik Duval. "Understanding Engagement with Interactive Public Displays: An Awareness Campaign in the Wild". In: *Proceedings of The International Symposium on Pervasive Displays*. PerDis '14. Copenhagen, Denmark: ACM, 2014, 180:180–180:185. ISBN: 978-1-4503-2952-1. DOI: 10.1145/2611009.2611020. URL: http://doi.acm.org/10.1145/2611009.2611020.

[Pel+18] Ella Peltonen, Eemil Lagerspetz, Jonatan Hamberg, Abhinav Mehrotra, Mirco Musolesi, Petteri Tapio Nurmi, and Sasu Tarkoma. "The Hidden Image of Mobile Apps: Geographic, Demographic, and Cultural Factors in Mobile Usage". In: *Proceedings of the 20th International Conference on Human-Computer Interaction with Mobile Devices and Services*. Apr. 2018. DOI: 10.1145/3229434.3229474. URL: https://www.cs.helsinki.fi/group/carat/pubs/mobilehci18.pdf.

[Pel+07] Peter Peltonen, Antti Salovaara, Giulio Jacucci, Tommi Ilmonen, Carmelo Ardito, Petri Saarikko, and Vikram Batra. "Extending Large-scale Event Participation with User-created Mobile Media on a Public Display". In: *Proceedings of the 6th International Conference on Mobile and Ubiquitous Multimedia*. MUM '07. Oulu, Finland: ACM, 2007, pp. 131–138. ISBN: 978-1-59593-916-6. DOI: 10.1145/1329469.1329487. URL: http://doi.acm.org/10.1145/1329469.1329487.

[Pet04] Eric Peterson. *Web Analytics Demystified: A Marketer's Guide to Understanding How Your Web Site Affects Your Business*. Celilo Group Media, 2004. ISBN: 0974358428. URL: https://www.amazon.com/Web-Analytics-Demystified-Marketers-Understanding/dp/0974358428?SubscriptionId=AKIAIOBINVZYXZQZ2U3A&tag=chimbori05-20&linkCode=xm2&camp=2025&creative=165953&creativeASIN=0974358428.

[PKK14] Thomas Phan, Swaroop Kalasapur, and Anugeetha Kunjithapatham. "Sensor fusion of physical and social data using Web SocialSense on smartphone mobile browsers". In: *2014 IEEE 11th Consumer Communications and Networking Conference (CCNC)*. Jan. 2014, pp. 98–104. DOI: 10.1109/CCNC.2014.6866555. URL: https://doi.org/10.1109/CCNC.2014.6866555.

[Qui] Quividi. *Is it possible to use the Quividi solution to change content depending on the profile and activity of the audience?* http://help.quividi.com/en/support/solutions/articles/6000038226-is-it-possible-to-use-the-quividi-solution-to-change-content-depending-on-the-profile-and-activity-of. (Visited on 10/22/2018).

[Qui16] Quividi. *Product Sheet: VidiReports*. Tech. rep. Quividi, Oct. 2016.

[RJK11] Harikrishna G.N. Rai, Kishore Jonna, and P. Radha Krishna. "Video Analytics Solution for Tracking Customer Locations in Retail Shopping Malls". In: *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. KDD '11. San Diego, California, USA: ACM, 2011, pp. 773–776. ISBN: 978-1-4503-0813-7. DOI: 10.1145/2020408.2020533. URL: http://doi.acm.org/10.1145/2020408.2020533.

[Ral+14] Swati Rallapalli, Aishwarya Ganesan, Krishna Chintalapudi, Venkat N. Padmanabhan, and Lili Qiu. "Enabling Physical Analytics in Retail Stores Using Smart Glasses". In: *Proceedings of the 20th Annual International Conference on Mobile Computing and Networking*. MobiCom '14. Maui, Hawaii, USA: ACM, 2014, pp. 115–126. ISBN: 978-1-4503-2783-1. DOI: 10.1145/2639108.2639126. URL: http://doi.acm.org/10.1145/2639108.2639126.

[RS13] Robert Ravnik and Franc Solina. “Audience Measurement of Digital Signage: Quantitative Study in Real-World Environment Using Computer Vision”. In: *Interacting with Computers* 25.3 (2013), p. 218. DOI: 10.1093/iwc/iws023. URL: http://dx.doi.org/10.1093/iwc/iws023.

[RSZ14] Robert Ravnik, Franc Solina, and Vesna Zabkar. “Modelling In-Store Consumer Behaviour Using Machine Learning and Digital Signage Audience Measurement Data”. In: *Video Analytics for Audience Measurement: First International Workshop, VAAM 2014, Stockholm, Sweden, August 24, 2014. Revised Selected Papers*. Ed. by Cosimo Distante, Sebastiano Battiato, and Andrea Cavallaro. Cham: Springer International Publishing, 2014, pp. 123–133. ISBN: 978-3-319-12811-5. DOI: 10.1007/978-3-319-12811-5_9. URL: http://dx.doi.org/10.1007/978-3-319-12811-5_9.

[Rem17] Remote Media Group Limited. *signagelive - Digital Signage Software Platform*. 2017. URL: https://signagelive.com (visited on 04/25/2017).

[RB05] Yvonne Rogers and Harry Brignull. “Subtle ice-breaking: encouraging socializing and interaction around a large public display”. In: *in Workshop on Public, Community. and Situated*. 2005. URL: http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.475.7738.

[RWC14] Nirupam Roy, He Wang, and Romit Roy Choudhury. “I Am a Smartphone and I Can Tell My User’s Walking Direction”. In: *Proceedings of the 12th Annual International Conference on Mobile Systems, Applications, and Services*. MobiSys ’14. Bretton Woods, New Hampshire, USA: ACM, 2014, pp. 329–342. ISBN: 978-1-4503-2793-0. DOI: 10.1145/2594368.2594392. URL: http://doi.acm.org/10.1145/2594368.2594392.

[RCS10] Antonio P. Russo, Salvador Anton Clave, and Noam Shoval. “Advanced Visitor Tracking Analysis in Practice: Explorations in the PortAventura Theme Park and Insights for a Future Research Agenda”. In: *Information and Communication Technologies in Tourism 2010: Proceedings of the International Conference in Lugano, Switzerland, February 10–12, 2010*. Ed. by Ulrike Gretzel, Rob Law, and Matthias Fuchs. Vienna: Springer Vienna, 2010, pp. 159–170. ISBN: 978-3-211-99407-8. DOI: 10.1007/978-3-211-99407-8_14. URL: http://dx.doi.org/10.1007/978-3-211-99407-8_14.

[RZ93] Roland T. Rust and Anthony J. Zahorik. “Customer satisfaction, customer retention, and market share”. In: *Journal of Retailing* 69.2 (1993), pp. 193–215. ISSN: 0022-4359. DOI: https://doi.org/10.1016/0022-4359(93)90003-2. URL: http://www.sciencedirect.com/science/article/pii/0022435993900032.

[SHT10] Javier San Agustin, John Paulin Hansen, and Martin Tall. “Gaze-based Interaction with Public Displays Using Off-the-shelf Components”. In: *Proceedings of the 12th ACM International Conference Adjunct Papers on Ubiquitous Computing - Adjunct*. UbiComp ’10 Adjunct. Copenhagen, Denmark: ACM, 2010, pp. 377–378. ISBN: 978-1-4503-0283-8. DOI: 10.1145/1864431.1864444. URL: http://doi.acm.org/10.1145/1864431.1864444.

[San+11] Addicam Sanjay, Shahzad Malik, Abhishek Ranjan, and Shweta Phadnis. *Metrics for Anonymous Video Analytics on Digital Signage*. White Paper. Intel Corporation, Aug. 2011.

[Sat+15] M. Satyanarayanan, P. Simoens, Y. Xiao, P. Pillai, Z. Chen, K. Ha, W. Hu, and B. Amos. “Edge Analytics in the Internet of Things”. In: *IEEE Pervasive Computing* 14.2 (Apr. 2015), pp. 24–31. ISSN: 1536-1268. DOI: 10.1109/MPRV.2015.32.

[Sca] Scala – Digital Signage Software. *Advanced Analytics – Measurable data for optimized messaging*. http://scala.com/advanced-analytics/. (Visited on 10/22/2018).

[SOC08] Jürgen Scheible, Timo Ojala, and Paul Coulton. "MobiToss: A Novel Gesture Based Interface for Creating and Sharing Mobile Multimedia Art on Large Public Displays". In: *Proceedings of the 16th ACM International Conference on Multimedia*. MM '08. Vancouver, British Columbia, Canada: ACM, 2008, pp. 957–960. ISBN: 978-1-60558-303-7. DOI: 10.1145/1459359.1459532. URL: http://doi.acm.org/10.1145/1459359.1459532.

[See] Seemetrix. *Seemetrix – Automated Audience Measurement*. https://seemetrix.net. (Visited on 10/22/2018).

[Sen+07] A. W. Senior, L. Brown, A. Hampapur, C. F. Shu, Y. Zhai, R. S. Feris, Y. L. Tian, S. Borger, and C. Carlson. "Video analytics for retail". In: *2007 IEEE Conference on Advanced Video and Signal Based Surveillance*. Sept. 2007, pp. 423–428. DOI: 10.1109/AVSS.2007.4425348.

[SSB06] Richard Sharp, James Scott, and Alastair R. Beresford. "Secure Mobile Computing Via Public Terminals". In: *Pervasive Computing*. Ed. by Kenneth P. Fishkin, Bernt Schiele, Paddy Nixon, and Aaron Quigley. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, pp. 238–253. ISBN: 978-3-540-33895-6.

[She+14] James She, Jon Crowcroft, Hao Fu, and Flora Li. "Convergence of Interactive Displays with Smart Mobile Devices for Effective Advertising: A Survey". In: *ACM Trans. Multimedia Comput. Commun. Appl.* 10.2 (Feb. 2014), 17:1–17:16. ISSN: 1551-6857. DOI: 10.1145/2557450. URL: http://doi.acm.org/10.1145/2557450.

[SI01] Garth B. D. Shoemaker and Kori M. Inkpen. "Single Display Privacyware: Augmenting Public Displays with Private Information". In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '01. Seattle, Washington, USA: ACM, 2001, pp. 522–529. ISBN: 1-58113-327-8. DOI: 10.1145/365024.365349. URL: http://doi.acm.org/10.1145/365024.365349.

[Sip+10] Andreas Sippl, Clemens Holzmann, Doris Zachhuber, and Alois Ferscha. "Real-Time Gaze Tracking for Public Displays". In: *Ambient Intelligence: First International Joint Conference, AmI 2010, Malaga, Spain, November 10-12, 2010. Proceedings*. Ed. by Boris de Ruyter, Reiner Wichert, David V. Keyson, Panos Markopoulos, Norbert Streitz, Monica Divitini, Nikolaos Georgantas, and Antonio Mana Gomez. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 167–176. ISBN: 978-3-642-16917-5. DOI: 10.1007/978-3-642-16917-5_17. URL: https://doi.org/10.1007/978-3-642-16917-5_17.

[Sla10a] Richard Slawsky. *Audience Measurement with Anonymous Video Analytics*. White Paper. DigitalSignageToday.com, 2010. URL: https://www.digitalsignagetoday.com/whitepapers/audience-measurement-with-anonymous-video-analytics/.

[Sla10b] Richard Slawsky. *Developments in Audience Measurement for Digital Signage*. White Paper. DigitalSignageToday.com, 2010. URL: https://www.digitalsignagetoday.com/whitepapers/developments-in-audience-measurement-for-digital-signage/.

[Sla16] Richard Slawsky. *Improving the Kiosk Experience with Analytics*. White Paper. KioskMarketplace.com, 2016. URL: https://www.kioskmarketplace.com/whitepapers/improving-the-kiosk-experience-with-analytics/.

[Sla11] Richard Slawsky. *Monetizing Anonymous Video Analytics*. White Paper. DigitalSignageToday.com, 2011. URL: https://www.digitalsignagetoday.com/whitepapers/monetizing-anonymous-video-analytics/.

[SBW12] Rajinder Sodhi, Hrvoje Benko, and Andrew Wilson. "LightGuide: Projected Visualizations for Hand Movement Guidance". In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '12. Austin, Texas, USA: ACM, 2012, pp. 179–188. ISBN: 978-1-4503-1015-4. DOI: 10.1145/2207676.2207702. URL: http://doi.acm.org/10.1145/2207676.2207702.

[Son] Sony Ziris™. http://pro.sony.com/bbsc/ssr/cat-digitalsignage/resource.solutions.bbsccms-assets-cat-digsignagedev-solutions-Ziris.shtml (Last accessed: October 2018).

[SAS11] Mani Srivastava, Tarek Abdelzaher, and Boleslaw Szymanski. "Human-centric sensing". In: *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* 370.1958 (2011), pp. 176–197. ISSN: 1364-503X. DOI: 10.1098/rsta.2011.0244. URL: http://rsta.royalsocietypublishing.org/content/370/1958/176.

[Sta98] William Stallings. "SNMP and SNMPv2: the infrastructure for network management". In: *IEEE Communications Magazine* 36.3 (Mar. 1998), pp. 37–43. ISSN: 0163-6804. DOI: 10.1109/35.663326. URL: https://doi.org/10.1109/35.663326.

[SFD06] Oliver Storz, Adrian Friday, and Nigel Davies. "Supporting content scheduling on situated public displays". In: *Computers & Graphics* 30.5 (2006), pp. 681–691. ISSN: 0097-8493. DOI: https://doi.org/10.1016/j.cag.2006.07.002. URL: http://www.sciencedirect.com/science/article/pii/S0097849306001191.

[Sto+06] Oliver Storz, Adrian Friday, Nigel Davies, Joe Finney, Corina Sas, and Jennifer Sheridan. "Public Ubiquitous Computing Systems: Lessons from the e-Campus Display Deployments". In: *IEEE Pervasive Computing* 5.3 (July 2006), pp. 40–47. ISSN: 1536-1268. DOI: 10.1109/MPRV.2006.56. URL: https://doi.org/10.1109/MPRV.2006.56.

[SM11] Martin Strohbach and Miquel Martin. "Toward a Platform for Pervasive Display Applications in Retail Environments". In: *IEEE Pervasive Computing* 10.2 (Apr. 2011), pp. 19–27. ISSN: 1536-1268. DOI: 10.1109/MPRV.2011.16. URL: https://doi.org/10.1109/MPRV.2011.16.

[TJS10] Thiago Teixeira, Deokwoo Jung, and Andreas Savvides. "Tasking Networked CCTV Cameras and Mobile Phones to Identify and Localize Multiple People". In: *Proceedings of the 12th ACM International Conference on Ubiquitous Computing*. UbiComp '10. Copenhagen, Denmark: ACM, 2010, pp. 213–222. ISBN: 978-1-60558-843-8. DOI: 10.1145/1864349.1864367. URL: http://doi.acm.org/10.1145/1864349.1864367.

[TP13] Omer Tene and Jules Polonetsky. "Big Data for All: Privacy and User Control in the Age ofAnalytics". In: *Northwestern Journal of Technology and Intellectual Property* 11.5 (2013).

[Tia+12] Phil Tian, Addicam V. Sanjay, Kunapareddy Chiranjeevi, and Shahzad Malik Malik. "Intelligent Advertising Framework for Digital Signage". In: *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. KDD '12. Beijing, China: ACM, 2012, pp. 1532–1535. ISBN: 978-1-4503-1462-6. DOI: 10.1145/2339530.2339773. URL: http://doi.acm.org/10.1145/2339530.2339773.

[Tom+14] Martin Tomitsch, Christopher Ackad, Oliver Dawson, Luke Hespanhol, and Judy Kay. "Who Cares About the Content? An Analysis of Playful Behaviour at a Public Display". In: *Proceedings of The International Symposium on Pervasive Displays*. PerDis '14. Copenhagen, Denmark: ACM, 2014, 160:160–160:165. ISBN: 978-1-4503-2952-1. DOI: 10.1145/2611009.2611016. URL: http://doi.acm.org/10.1145/2611009.2611016.

[VO11] Ville Valkama and Timo Ojala. "Stakeholder Value Propositions on Open Community Testbed of Interactive Public Displays". In: *Proceedings of the International Conference on Management of Emergent Digital EcoSystems*. MEDES '11. San Francisco, California: ACM, 2011, pp. 107–113. ISBN: 978-1-4503-1047-5. DOI: 10.1145/2077489.2077509. URL: http://doi.acm.org/10.1145/2077489.2077509.

[Ven+07] P. L. Venetianer, Z. Zhang, A. Scanlon, Y. Hu, and A. J. Lipton. “Video verification of point of sale transactions”. In: *2007 IEEE Conference on Advanced Video and Signal Based Surveillance*. June 2007, pp. 411–416. DOI: 10.1109/AVSS.2007.4425346. URL: https://doi.org/10.1109/AVSS.2007.4425346.

[VB04] Daniel Vogel and Ravin Balakrishnan. “Interactive Public Ambient Displays: Transitioning from Implicit to Explicit, Public to Personal, Interaction with Multiple Users”. In: *Proceedings of the 17th Annual ACM Symposium on User Interface Software and Technology*. UIST ’04. Santa Fe, NM, USA: ACM, 2004, pp. 137–146. ISBN: 1-58113-957-8. DOI: 10.1145/1029632.1029656. URL: http://doi.acm.org/10.1145/1029632.1029656.

[WW95] C. A. Waldspurger and E. Weihl. W. *Stride Scheduling: Deterministic Proportional-Share Resource Management*. Tech. rep. Cambridge, MA, USA, 1995.

[WW94] Carl A. Waldspurger and William E. Weihl. “Lottery Scheduling: Flexible Proportional-share Resource Management”. In: *Proceedings of the 1st USENIX Conference on Operating Systems Design and Implementation*. OSDI ’94. Monterey, California: USENIX Association, 1994. URL: http://dl.acm.org/citation.cfm?id=1267638.1267639.

[WBM13] Robert Walter, Gilles Bailly, and Jörg Müller. “StrikeAPose: Revealing Mid-air Gestures on Public Displays”. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI ’13. Paris, France: ACM, 2013, pp. 841–850. ISBN: 978-1-4503-1899-0. DOI: 10.1145/2470654.2470774. URL: http://doi.acm.org/10.1145/2470654.2470774.

[WBG12] Miaosen Wang, Sebastian Boring, and Saul Greenberg. “Proxemic Peddler: A Public Advertising Display That Captures and Preserves the Attention of a Passerby”. In: *Proceedings of the 2012 International Symposium on Pervasive Displays*. PerDis ’12. Porto, Portugal: ACM, 2012, 3:1–3:6. ISBN: 978-1-4503-1414-5. DOI: 10.1145/2307798.2307801. URL: http://doi.acm.org/10.1145/2307798.2307801.

[Wan+92] Roy Want, Andy Hopper, Veronica Falcão, and Jonathan Gibbons. “The Active Badge Location System”. In: *ACM Trans. Inf. Syst.* 10.1 (Jan. 1992), pp. 91–102. ISSN: 1046-8188. DOI: 10.1145/128756.128759. URL: http://doi.acm.org/10.1145/128756.128759.

[WW14] Julie R. Williamson and John Williamson. “Analysing Pedestrian Traffic Around Public Displays”. In: *Proceedings of The International Symposium on Pervasive Displays*. PerDis ’14. Copenhagen, Denmark: ACM, 2014, 13:13–13:18. ISBN: 978-1-4503-2952-1. DOI: 10.1145/2611009.2611022. URL: http://doi.acm.org/10.1145/2611009.2611022.

[WH08] Oliver Woodman and Robert Harle. “Pedestrian Localisation for Indoor Environments”. In: *Proceedings of the 10th International Conference on Ubiquitous Computing*. UbiComp ’08. Seoul, Korea: ACM, 2008, pp. 114–123. ISBN: 978-1-60558-136-1. DOI: 10.1145/1409635.1409651. URL: http://doi.acm.org/10.1145/1409635.1409651.

[Yan+07] T. Yang, F. Chen, D. Kimber, and J. Vaughan. “Robust People Detection and Tracking in a Multi-Camera Indoor Visual Surveillance System”. In: *2007 IEEE International Conference on Multimedia and Expo*. July 2007, pp. 675–678. DOI: 10.1109/ICME.2007.4284740.

[Zha+10] Peng Zhang, Tony Thomas, Sabu Emmanuel, and Mohan S. Kankanhalli. “Privacy Preserving Video Surveillance Using Pedestrian Tracking Mechanism”. In: *Proceedings of the 2Nd ACM Workshop on Multimedia in Forensics, Security and Intelligence*. MiFor ’10. Firenze, Italy: ACM, 2010, pp. 31–36. ISBN: 978-1-4503-0157-2. DOI: 10.1145/1877972.1877983. URL: http://doi.acm.org/10.1145/1877972.1877983.

[ZBG13] Yanxia Zhang, Andreas Bulling, and Hans Gellersen. "SideWays: A Gaze Interface for Spontaneous Interaction with Situated Displays". In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '13. Paris, France: ACM, 2013, pp. 851–860. ISBN: 978-1-4503-1899-0. DOI: 10.1145/2470654.2470775. URL: http://doi.acm.org/10.1145/2470654.2470775.

[Zho+08] Maohua Zhong, Congling Shi, Xuwei Tu, Tairan Fu, and Li He. "Study of the human evacuation simulation of metro fire safety analysis in China". In: *Journal of Loss Prevention in the Process Industries* 21.3 (2008), pp. 287–298. DOI: 10.1016/j.jlp.2007.08.001. URL: https://www.sciencedirect.com/science/article/pii/S0950423007001209.