On User Privacy for Location-based Services

Anand S. Gajparia

Technical Report
RHUL–MA–2007–7
6 June 2007

Department of Mathematics
Royal Holloway, University of London
Egham, Surrey TW20 0EX, England
http://www.rhul.ac.uk/mathematics/techreports

On User Privacy for Location-based Services

Anand S. Gajparia

Thesis submitted to the University of London
for the degree of Doctor of Philosophy

Information Security Group
Department of Mathematics
Royal Holloway, University of London
2007

Declaration

These doctoral studies were conducted under the supervision of Chris J. Mitchell and Peter Wild. The work presented in this thesis is the result of original research carried out by myself, in collaboration with others, whilst enrolled in the Department of Mathematics as a candidate for the degree of Doctor of Philosophy. This work has not been submitted for any other degree or award in any other university or educational establishment.

Anand Gajparia

Acknowledgements

I would like to thank my supervisor, Professor C. J. Mitchell for his invaluable support and guidance. I would also like to thank Dr Chan Yeob Yeun for his constructive criticism and feedback.

Toshiba Telecommunications Laboratory provided me with sponsorship and resources without which this thesis would not have been possible. I am eternally grateful.

Finally, I would like to thank my parents, Jaimi, and my other friends and family, too many to mention, for all their help and encouragement.

Abstract

This thesis investigates user privacy concerns associated with the use of location based services. We begin by introducing various privacy schemes relevant to the use of location based services.

We introduce the notion of constraints, i.e. statements limiting the use and distribution of Location Information (LI), i.e. data providing information regarding a subject's location. Constraints can be securely bound to LI, and are designed to reduce threats to privacy by controlling its dissemination and use. The various types of constraint which may be required are also considered. The issues and risks with the possible use of constraints are discussed, as are possible solutions to these hazards.

To address some of the problems that have been identified with the use of constraints, we introduce the notion of an LI Preference Authority (LIPA). A LIPA is a trusted party which can examine LI constraints and make decisions about LI distribution without revealing the constraints to the entity requesting the LI. This is achieved by encrypting both the LI and the constraints with a LIPA encryption key, ensuring that the LI is only revealed at the discretion of the LIPA.

We further show how trusted computing can be used to enhance privacy for LI. We focus on how the mechanisms in the Trusted Computing Group specifications can be used to enable the holder of LI to verify the trustworthiness of a remote host before transferring the LI to that remote device. This provides greater assurance to end users that their expressed preferences for the handling of personal information will be respected.

The model for the control of LI described in this thesis has close parallels to models controlling the dissemination and use of other personal information. In particular, Park and Sandhu have developed a general access control model intended to address issues such as Digital Rights Management, code authorisation, and the control of personal data. We show how our model for LI control fits into this general access control model.

We present a generic service which allows a device to discover the location of other

devices in ad hoc networks. The advantages of the service are discussed in several scenarios, where the reliance on an infrastructure such as GPS satellites or GSM cellular base stations is not needed. An outline of the technology which will be needed to realise the service is given, along with a look at the security issues which surround the use of this location discovery service.

Finally, we provide conclusions and suggestions for future work.

Contents

1 Introduction **15**
1.1 Motivation and Challenges . . . 15
1.1.1 The P3P Framework . . . 15
1.1.2 Behavioural and Characteristic Information . . . 18
1.1.3 Location Information — Our Definition . . . 19
1.1.4 The Nature and Use of LI . . . 19
1.1.5 LI Privacy . . . 21
1.2 Structure of Thesis . . . 22
1.3 Contributions of Thesis . . . 23
1.4 Publications . . . 23

2 Notation and Cryptographic Primitives **25**
2.1 Introduction . . . 26
2.2 Security Services . . . 26
2.2.1 Confidentiality . . . 26
2.2.2 Integrity . . . 26
2.2.3 Access Control . . . 26
2.2.4 Authentication . . . 27
2.2.5 Non-Repudiation . . . 27
2.2.6 Privacy . . . 27
2.2.7 Anonymity . . . 27
2.2.8 Pseudonym . . . 28
2.3 Security Mechanisms . . . 28
2.3.1 Secret Key (Symmetric) Encryption . . . 28
2.3.2 Public Key (Asymmetric) Encryption . . . 29
2.3.3 Cryptographic Hash Functions . . . 30
2.3.4 Digital Signatures . . . 30
2.3.5 Time Stamps . . . 31
2.4 Security and Privacy . . . 31

3 Privacy and Location-Based Services **33**
3.1 Introduction . . . 34
3.2 Threats to LI . . . 36
3.2.1 Unauthorised LI Use . . . 36
3.2.2 Eavesdropping Threat . . . 37
3.2.3 Unauthorised LI Modification . . . 37
3.3 Using Access Constraints . . . 38
3.4 Information Reduction . . . 40

3.5 Using Regulation, Standards, and Guidelines 43
3.5.1 Regulatory Bodies Providing User Location Privacy 43
3.5.2 Self Regulation 47
3.5.3 Standards and Guidelines 48
3.6 Summary 55

4 Constraints and Location Based Services 56
4.1 Introduction 57
4.2 A Model for the Use of LI 58
4.2.1 The Roles 59
4.3 Using Constraints with LI 61
4.3.1 Constraint Types 62
4.3.2 Uses of Constraints 64
4.4 Limitations of Constraints 66
4.4.1 Difficulties in Preventing and Detecting Constraint Abuse . . 66
4.4.2 LI Constraint Predicaments 67
4.5 Combining Constraints with Auditability 68
4.6 Specification of Constraints 70
4.7 Summary 71

5 The Location Information Preference Authority 72
5.1 Introduction 73
5.2 A Mechanism to Provide Security for Constraints 74
5.2.1 Overview of the Mechanism 74
5.2.2 Requirements for Use of the Mechanism 77
5.2.3 LI Token Creation 78
5.2.4 LI Distribution 79
5.2.5 LI Use 80
5.3 Billing 82
5.4 Performance Analysis 83
5.4.1 Assumptions 84
5.4.2 Storage Requirements 85
5.4.3 Message Exchanges 86
5.4.4 Operations Performed 87
5.5 Security Analysis 89
5.6 Summary 91

6 Using Trusted Computing to Enhance Location Privacy 92
6.1 Introduction 93
6.2 Trusted Computing 94
6.3 Next-Generation Secure Computing Base 96
6.4 Personal Information Model 97
6.5 Scenarios 98
6.5.1 Registration Scenario 98
6.5.2 Location-Based Service Scenario 99
6.5.3 Medical Records 99
6.6 TCG Mechanisms 100

6.6.1 TPM Identities 100
6.6.2 TCG Measuring, Reporting and Storing Processes 102
6.6.3 Sealing Data 103
6.7 Protecting Personal Information Using Trusted Computing 104
6.7.1 Overview 104
6.7.2 Practicality 106
6.7.3 Using Trusting Computing with PI 107
6.7.4 Constraints, LIPA and LI Tokens 109
6.8 Other Approaches to Privacy Protection Using Trusted Computing . 110
6.9 Analysis 110
6.10 Conclusions 111

7 Location-Based Services and the Usage Control Model 112
7.1 Introduction 113
7.2 LI Entities 115
7.3 Constraints 115
7.3.1 LI Gatherer Constraints 116
7.3.2 LI Distribution Constraints 116
7.3.3 LI Use Constraints 118
7.3.4 Difficulties in Implementing Constraints 118
7.4 An Introduction to UCON 119
7.5 Modeling LI Constraints in UCON_{ABC} 121
7.5.1 Authorisation Restrictions 123
7.5.2 Conditional Restrictions 123
7.5.3 Model Definitions 124
7.5.4 Modelling LI Constraints in the UCON Model 126
7.6 Analysis, Future Work and Conclusions 129

8 Generating LI Using Ad Hoc Networks 131
8.1 Introduction 132
8.2 Terminology 134
8.3 Ad Hoc Networks 135
8.4 The Location Discovery Service 136
8.5 Requirements and Architecture 137
8.5.1 Requirements 137
8.5.2 Infrastructure Based Tracking 138
8.5.3 Ad Hoc Tracking 139
8.5.4 Other Scenarios 140
8.6 Location Technology Overview 140
8.6.1 The GPS Method 140
8.6.2 The Smart Antenna Method 141
8.6.3 DOA for Omnidirectional Antennas 142
8.7 Security Requirements 142
8.8 Security Solutions 144
8.9 Related Work 146
8.10 Future Work 150
8.11 Summary and Conclusions 151

9 Conclusions **152**
9.1 Summary and Conclusions 152
9.2 Suggestions for Further Work 153

Bibliography **155**

A Information from web-pages **171**
A.1 Introduction 171
A.2 Cingular Wireless Privacy Policy 171
A.3 Features and Services Information for Former AT&T Wireless Users - mMode 188

List of Figures

4.1 LI usage tree 66

5.1 LIPA summary 75
5.2 LI-related transmission for a mobile UD 76

7.1 The stages at which control can be applied to LI 115

List of Tables

5.1 LI Lengths 84
5.2 Message Exchanges 87
5.3 Storage Requirements 87
5.4 LI Token Generation Operations 88
5.5 LIPA Operations When LI Token is Received 88
5.6 LIPA Operations When LI is Sent 88
5.7 LBS Provider Operations 89

7.1 Instances of UCON_{ABC} 121
7.2 Mapping LI constraint types to UCON attribute types 122

Abbreviations

AAA:	Authorisation, Authentication and Accounting
AES:	Advanced Encryption Standard
AIK:	Attestation Identity Key
AODV:	Ad hoc On-demand Distance Vector
BITS:	Boot Integrity Token System
CA:	Certification Authority
CPNI:	Customer Proprietary Network Information
CRTM:	Core Root of Trust for Measurement
CTIA:	Cellular Telecommunications and Internet Association
DAA:	Direct Anonymous Attestation
DES:	Data Encryption Standard
DHCP:	Dynamic Host Configuration Protocol
DMCA:	Digital Millennium Copyright Act
DOA:	Direction Of Arrival
DRM:	Digital Rights Management
DSA:	Digital Signature Algorithm
DSR:	Dynamic Source Routing
EOTD:	Enhanced Observed Time Difference
EU:	European Union
FCC:	Federal Communications Commission
Geopriv:	Geographic Location/Privacy
GML:	Geography Markup Language
GPS	Global Positioning System
GSM:	Global System for Mobile communications
ID:	Identity
IP:	Internet Protocol
IPSec:	Internet Protocol Security
IPv6:	Internet Protocol Version 6

IETF:	Internet Engineering Task Force
IMSI:	International Mobile Subscriber Identifier
KDC:	Key Distribution Centre
LAN:	Local Area Network
LAR:	Location Aided Routing
LBS:	Location Based Service
LG:	Location Gatherer
LI:	Location Information
LIPA:	Location Information Preference Authority
LO:	Location Object
LOCI:	Location Information
LR:	Location Recipient
LS:	Location Server
MAC Address:	Media Access Control Address
MD4:	Message Digest Four
MD5:	Message Digest Five
NGSCB:	Next-Generation Secure Computing Base
NIST:	National Institute for Standards and Technology
OAEP:	Optimal Asymmetric Encryption Padding
OECD:	Organisation for Economic Cooperation and Development
P3P:	Platform for Privacy Preferences Project
PET:	Privacy Enhancing Technology
PBX:	Public Branch eXchange
PCR:	Platform Configuration Register
PDA:	Personal Digital Assistant
PI:	Personal Information
PIDF:	Presence Information Data Format
PIDF-LO:	Presence Information Data Format Location Object
PIR:	Private Information Retrieval
PKI:	Public Key Infrastructure
RADAR:	Radio Detection And Ranging
RADIUS:	Remote Authentication Dial-In User Service
RFC:	Request For Comments
RFID:	Radio Frequency Identification
RH:	Rule Holder

RSA:	Rivest-Shamir-Adleman
RTM:	Root of Trust for Measuring Integrity Metrics
SAML:	Security Assertion Markup Language
SDSI:	Simple Distributed Security Infrastructure
SHA-1:	Secure Hash Algorithm revision One
S/MIME:	Secure/Multipurpose Internet Mail Extensions
SP:	Service Provider
SPKI:	Simple Public Key Infrastructure
SSO:	Single Sign On
TA:	Timing Advance
TC:	Trusted Computing
TCG:	Trusted Computing Group
TCP:	Trusted Computing Platform
TCPA:	Trusted Computing Platform Alliance
TGT:	Ticket Granting Ticket
TP:	Trusted Platform
TPM:	Trusted Platform Module
TPS:	Trusted Platform Subsystem
TSS:	Trusted Software Stack
UCON:	Usage Control
UD:	User Device
URL:	Uniform Resource Locator
UTM:	Universal Transverse Mercator
US:	United States
VOIP:	Voice Over Internet Protocol
VPN:	Virtual Private Network
W3C:	World Wide Web Consortium
WAP:	Wireless Access Protocol
WLAN:	Wireless Local Area Network
XML:	Extensible Markup Language

Chapter 1

Introduction

Contents

1.1 Motivation and Challenges 15
1.1.1 The P3P Framework 15
1.1.2 Behavioural and Characteristic Information 18
1.1.3 Location Information — Our Definition 19
1.1.4 The Nature and Use of LI 19
1.1.5 LI Privacy 21
1.2 Structure of Thesis 22
1.3 Contributions of Thesis 23
1.4 Publications 23

This chapter introduces the thesis. It gives an outline of the motivation for the research and the major challenges that it addresses. It also describes the structure of the thesis and its main contributions.

1.1 Motivation and Challenges

A number of data-types can be classified as personal information. This section discusses some of these, and also describes where Location Information fits.

1.1.1 The P3P Framework

The World-Wide Web Consortium (W3C) Platform for Privacy Preferences Project (P3P) provides a framework for expressing a user's requirements for privacy when accessing Web sites, to enable a user system to learn about that Web site's privacy

practices. A computer agent can then compare the privacy requirements of the user to those of the Web site being visited, and advise the user accordingly. Currently, P3P defines 16 data categories [31]. These are described below.

- **Physical Contact Information.** Information that allows an individual to be contacted or located in the physical world — such as a telephone number or street address.
- **Online Contact Information.** Information that allows an individual to be contacted or located on the Internet — such as an email address. Often, this information is independent of the specific computer used to access the network.
- **Unique Identifiers.** Non-financial identifiers, excluding government-issued identifiers, issued for the purpose of consistently identifying or recognising the individual. These include identifiers issued by a Web site or service.
- **Purchase Information.** Information actively generated by the purchase of a product or service, including information about the method of payment.
- **Financial Information.** Information about an individual's finances including account status and activity information such as account balance, payment or overdraft history, and information about an individual's purchases or use of financial instruments including credit or debit card information. Information which is just derived from a discrete purchase by an individual, as described in "Purchase Information," does not come under the definition of "Financial Information."
- **Computer Information.** Information about the computer system that the individual is using to access the network — such as the IP address, domain name, browser type or operating system.
- **Navigation and Click-stream Data.** Data passively generated by browsing the Web site — such as which pages are visited, and how long users stay on each page.
- **Interactive Data.** Data actively generated from, or reflecting explicit interactions with, a service provider through its site — such as queries to a search engine, or logs of account activity.

- **Demographic and Socioeconomic Data.** Data about an individual's characteristics — such as gender, age, and income.
- **Content.** The words and expressions contained in the body of a communication — such as the text of an email, bulletin board postings, or chat room communications.
- **State Management Mechanisms.** Mechanisms for maintaining a stateful session with a user, or automatically recognising users who have visited a particular site or accessed particular content previously — such as HTTP cookies.
- **Political Information.** Membership in, or affiliation with, groups such as religious organisations, trade unions, professional associations, political parties, etc.
- **Health Information.** Information about an individual's physical or mental health, sexual orientation, use of or inquiry into health care services or products, and purchase of health care services or products.
- **Preference Data.** Data about an individual's likes and dislikes — such as favorite colour or musical tastes.
- **Location Data.** Information that can be used to identify an individual's current physical location and track them as their location changes — such as GPS position data.
- **Government-issued Identifiers.** Identifiers issued by a government for the purpose of consistently identifying the individual.
- **Other.** Other types of data not captured by the above definitions.

Location Information, as described in this thesis, extends that which is described as 'Location Data', in P3P. In addition to physically locating an individual, we also include information that can allow an entity to infer an individual's location. For example, previous behaviour may indicate where an individual is likely to be located in the future. That is, as we use the term in this thesis, Location Information might include data derived from a large number of information types, not just that which is categorised by P3P as Location Data.

1.1.2 Behavioural and Characteristic Information

The term 'behavioural information' refers to an attribute of an individual that can change depending on one or more other factors. The information and the factor affecting the change may be described as behavioural information. For example, a person may be at particular locations depending on the time of day. Predicting the behaviour of an individual can rely on a certain amount of behavioural information.

Behavioural information differs from characteristic information, which describes information about an individual that is unlikely to change, such as his/her name, or date of birth.

With this in mind, characteristic information includes the following categories of personal information from the list above: Physical Contact Information, Online Contact Information, Unique Identifiers, Demographic and Socioeconomic Data, Political Information, Health Information, and Government-issued Identifiers.

Behavioural information includes Location Data, Preference Data, Purchase Information, Financial Information, Computer Information, Navigation and Click-stream Data, Interactive Data, Content, and State Management Mechanisms. Although computer information may be static, it may also be possible to use this information to infer that an individual is located at a certain computer at a given time.

Controlling the dissemination and use of behavioural information requires more information about the current state of an individual than the control of characteristic information. For example, an individual may not want their location to be known at certain times, or may not want their location divulged when they are at a given place. Controlling characteristic information may require only a yes or no statement when it is being distributed to certain entities. An exception to this may be when a user requires their data to be manipulated before it is sent. For example, when their health information is being passed on for research purposes, they may not want to be identified.

1.1.3 Location Information — Our Definition

This thesis investigates a certain type of behavioural information, which we refer to as Location Information.

- **Location Information (LI).** This is data which provides information regarding an LI subject's location. LI may occur in many forms. In general, we can divide LI into two types, namely *Inferred* LI and *Actual* LI.
  - **Actual LI** refers to a directly calculated geographical location. This type of data indicates, to some degree of accuracy, the physical location of an LI subject.
  - **Inferred LI** is, by contrast, obtained by implication. For example, if a user is present on a network, this may imply that he or she is likely to be within a certain vicinity, although no specific calculation of geographical LI has taken place.

Actual LI is usually generated by a specialist entity referred to in this thesis as an LI gatherer — see Chapter 4 for further information. Another entity that plays an important role in this model is the entity about which LI is being gathered. This is described below.

- **LI subject.** An LI subject is the entity about which LI is being gathered, managed and used. This entity is most commonly a device owned by a user.

1.1.4 The Nature and Use of LI

Location Information differs from other types of behavioural data for a number of reasons. First, methods for generating this data may vary, i.e., LI data may be gathered on a mobile device, or it may be gathered by a third party. This adds to the challenge of expressing preferences about how an LI subject may want this data to be handled.

LI by itself is simply data. LI typically takes the form of co-ordinates indicating a geographical location. Context is added to this when it is used together with a map, or in the provision of a location-based service. Location-based services allow this data to be used as the basis of service provision. There are numerous motivations for the use of location-based services by a user. Examples includes the provision of emergency assistance, routing assistance, and the location of goods for logistics providers.

As devices used for wireless communication become increasingly ubiquitous and mobile, it is becoming apparent that location-based services will play an important role in the evolution of ambient networking. Location-based services use LI to allow an LI subject, i.e. the individual to whom the LI applies, or some other entity, to exploit this information to support the provision of one or more services. These range from allowing an emergency service to locate an LI subject, as is the case with E911 in the United States of America [50], to an authentication service based on the location of an LI subject [36]. Location-based media, another type of location-based service, where multi-media is delivered to a mobile device depending on its location, is another noteworthy example. An example of the use of location-based media is provided by History Unwired [48]. This allows a visitor to Venice to receive multi-media information about local culture and history on their mobile device based on their location. Services that allow members of a social community to interact based on their location have also been identified. Such services may allow a user to chat with other community members that are located nearby. An example of such a service is FriendZone [20]. Location-based services are also likely to play a significant role as a vehicular technology [129], including the support of navigational services and road toll schemes.

For the successful deployment of location-based services, various challenges need to be addressed. These challenges include interoperability with existing infrastructure, accuracy and user privacy concerns. Privacy of user location information (LI) is the main focus of this thesis.

LI can be generated in many ways, including using various forms of Global Positioning System (GPS) and Enhanced Observed Time Difference (EOTD) technologies [2]. These technologies are commonly implemented in mobile devices, and

improvements in such techniques have been motivated by E911 [79]. GPS uses 24 satellites to enable the calculation of the LI of a subject using triangulation techniques. EOTD calculates LI by observing time differences in transmissions between a user device and a base station. The RADAR [9] system determines the location of devices by calculations using signal strength and signal-to-noise ratios in IEEE 802.11 networks. Other technologies used to determine location include those using fixed receivers that detect when a transmitter comes within their proximity. Examples of such systems include Active Badge [149], that uses diffuse infrared technology, and Active Bat [66] and Cricket [113], that use ultrasonic emissions. For further information about location systems, the reader is referred to [79] (we also provide a brief review of such systems in chapter 8).

Location Information, although commonly gathered by a third party, may also be gathered by the individual. For example, an individual may gather their own location data using GPS and pass it to a service provider, or a base station may do this for them.

1.1.5 LI Privacy

Unfortunately, LI may also be used for malicious purposes. Obvious privacy violations include when LI is passed to entities that the LI subject does not wish to have this information. This could arise when a permitted entity passes LI to an unauthorised entity. By divulging a series of locations, unauthorised access to LI could enable a malicious party to track an LI subject, e.g. to stalk the LI subject. If a malicious party was able to find out that an anonymous message was from a certain location, and that a particular LI subject was regularly at that location at the time that the message was sent, then they could deduce that the message was sent by that LI subject.

Another undesirable use for LI is location-based spam [32]. This refers to unsolicited messages sent to a device based on its location [83]. Another scenario where a privacy violation may take place is when information can be deduced through LI. For example, it may be possible to deduce the place of residence and work of an LI subject by identifying where they are located during the night and during the day.

Intuitively, privacy of LI can only be gained by limiting its distribution. This would mean that only those authorised by the LI subject would be able to gain possession of LI. Securing the privacy of LI is an issue which needs to be addressed in order to gain the trust of consumers for such services.

This thesis describes an architecture for the generation, distribution and use of LI. The main aim of this architecture is to provide an infrastructure to support the provision of privacy for users of LI.

1.2 Structure of Thesis

The remainder of this thesis is divided into 8 chapters. An overview of each chapter is given below.

- **Chapter 2** is a preliminary chapter that introduces the notation and the cryptographic primitives used throughout the remainder of the thesis.
- **Chapter 3** provides a review of current research and technologies in user privacy and location-based services.
- **Chapter 4** describes how constraints can be used to extend user control over LI, even after it has been distributed. The types of constraint required for such control are considered, as well as issues and risks with the possible use of constraints. Possible solutions to these hazards are also discussed.
- **Chapter 5** introduces the notion of a LI Preference Authority (LIPA). This is a trusted party which can examine LI constraints and make decisions about LI distribution without revealing the constraints to the entity requesting the LI. We look at how the LIPA can be used to resolve some of the risks involved with the use of constraints.
- **Chapter 6** describes the use of trusted computing to enhance location privacy.
- **Chapter 7** investigates access control for LI in the context of the UCON_{ABC} model. The specification of the architecture developed in chapters 4–6 using the UCON_{ABC} model is explored. This provides further insights into the

operation of the architecture, as well as providing a use case for the UCON_{ABC} model.

- **Chapter 8** describes a generic location service for use in ad hoc networks. It describes requirements on the technology needed to achieve the service, and fundamental security issues which surround the use of such a service.
- Finally, **Chapter 9** summarises the contents of this thesis, and gives conclusions. It also suggests directions for future research.

1.3 Contributions of Thesis

This thesis proposes a number of ways of enhancing the privacy of user location information in the context of location-based services. These are as follows:

- A method, in the form of constraints, for users to control LI after it has been distributed;
- A scheme, using these constraints, to distribute LI without unnecessarily divulging these constraints to other parties;
- A scheme, using trusted computing, to ensure the proper handling of LI by third-parties;
- An investigation of access control for LI using the UCON_{ABC} model;
- An investigation into the security requirements for a location service in an ad hoc network.

1.4 Publications

Publications describing some of the research results in this thesis are listed below.

- A. S. Gajparia. On location-based services and the usage control model (extended abstract). In *Western European Workshop on Research in Cryptology*, pages 74–77. WEWoRC Conference Records, Leuven, Belgium, July 2005.

- A. S. Gajparia and C. J. Mitchell. Enhancing user privacy using trusted computing. In C. J. Mitchell, editor, *Trusted Computing*, chapter 8, pages 239–249. IEE, Hertfordshire, UK, 2005.

- A. S. Gajparia, C. J. Mitchell, and C. Y. Yeun. Using constraints to protect personal location information. In *Proceedings of the 58th IEEE Vehicular Technology Conference (VTC 2003-Fall)*, volume 3, pages 2112–2116. IEEE Press, Piscataway, NJ, USA, October 2003.

- A. S. Gajparia, C. J. Mitchell, and C. Y. Yeun. The location information preference authority: Supporting user privacy in location based services. In S. Liimatainen and T. Virtanen, editors, *Proceedings of Nordsec 2004, the 9th Nordic Workshop on Secure IT systems*, pages 91–96. Helsinki University of Technology, Finland, November 2004.

- A. S. Gajparia, C. J. Mitchell, and C. Y. Yeun. Supporting user privacy in location based services. *IEICE Transactions*, E88-B(7):2848–2855, July 2005.

Chapter 2

Notation and Cryptographic Primitives

Contents

2.1 Introduction **26**
2.2 Security Services **26**
2.2.1 Confidentiality 26
2.2.2 Integrity 26
2.2.3 Access Control 26
2.2.4 Authentication 27
2.2.5 Non-Repudiation 27
2.2.6 Privacy 27
2.2.7 Anonymity 27
2.2.8 Pseudonym 28
2.3 Security Mechanisms **28**
2.3.1 Secret Key (Symmetric) Encryption 28
2.3.2 Public Key (Asymmetric) Encryption 29
2.3.3 Cryptographic Hash Functions 30
2.3.4 Digital Signatures 30
2.3.5 Time Stamps 31
2.4 Security and Privacy **31**

This chapter defines security terms used throughout this thesis. The chapter begins with an introduction in section 2.1. Definitions for security services are provided in section 2.2, and definitions for security mechanisms are given in section 2.3.

2.1 Introduction

In this chapter we provide definitions of fundamental security notions which will be used throughout the remainder of the thesis.

Many of the definitions have been taken from ISO 7498-2 [76]. The definitions for anonymity and pseudonym were extracted from [111]. Various books [37, 76, 95] have also been used to derive these definitions.

2.2 Security Services

In the context of computer communications, the main security safeguards are known as security services [52]. The services discussed in this thesis are defined below.

2.2.1 Confidentiality

Confidentiality is the property that information is not made available or disclosed to unauthorised individuals, entities, or processes [76].

2.2.2 Integrity

Integrity is the property that data has not been altered or destroyed in an unauthorised manner [76].

2.2.3 Access Control

Access control is the prevention of unauthorised use of a resource, including the prevention of use of a resource in an unauthorised manner [76].

There are two fundamental ways of specifying an access control policy [62]. These are known as discretionary access control and mandatory access control. Discretionary

access control allows the owner of a resource to specify entities that are allowed access to a resource. Mandatory access control describes the case where access to a resource is determined by a system-wide policy.

2.2.4 Authentication

Authentication provides assurance of the identity of an entity. This entity is either the source of certain information, in which case the service is known as data origin authentication, or a communicating entity (a person or system), in which case the service is known as entity authentication [95].

2.2.5 Non-Repudiation

Non-repudiation protects against one party to a communication exchange later falsely denying that the exchange occurred [52].

2.2.6 Privacy

Privacy is the right of individuals to control or influence what information related to them may be collected and stored and by whom, and to whom that information may be disclosed [76].

2.2.7 Anonymity

Anonymity refers to a situation in which the author of an action is not identifiable within a set of subjects, known as the anonymity set [111]. The anonymity set is the set of all possible subjects which might have caused the action.

2.2.8 Pseudonym

Pseudonyms are identifiers of subjects [111]. The subject that may be identified by the pseudonym is the holder of the pseudonym. Unique pseudonyms can be used to realise accountability. Some possible properties of pseudonyms are listed below.

- **Public pseudonym.** In this case, the link between the pseudonym and a subject is publicly known from the point at which the pseudonym is first defined.
- **Initially non-public pseudonym.** Here the link between an initially non-public pseudonym and a subject may be known by certain parties, but is not public, at least initially.
- **Initially unlinkable pseudonym.** In this case, the link between an initially unlinkable pseudonym and a subject is, at least initially, not known to anybody with the possible exception of the holder himself/herself.

2.3 Security Mechanisms

Security mechanisms are used to realise the security services discussed in section 2.2. We now briefly review those mechanisms of particular importance to this thesis.

2.3.1 Secret Key (Symmetric) Encryption

Secret key (symmetric) encryption mechanisms can be used to provide confidentiality services. Let $\{E_e : e \in K\}$ be a set of encryption transformations, and $\{D_d : d \in K\}$ be the corresponding decryption transformations, where K is the key space. For each associated encryption/decryption key pair (e, d), it is computationally easy to determine d from e, and e from d. The use of symmetric encryption relies on the encrypting and decrypting parties sharing a secret key.

Examples of secret key encryption schemes include the Advanced Encryption Stan-

dard (AES) [99], DES and RC5.

2.3.2 Public Key (Asymmetric) Encryption

Public key (asymmetric) encryption mechanisms can also be used to provide confidentiality services. Let $\{E_e : e \in K\}$ be a set of encryption transformations and $\{D_d : d \in K\}$ be the corresponding decryption transformations, where K is the key space. For any pair of encryption/decryption keys, (e, d), the pair has the property that, given knowledge of e, it is computationally infeasible, given random ciphertext $c \in C$, to find a message $m \in M$ for which $E_e(m) = c$. This also implies that, given e, it is infeasible to determine d. Asymmetric encryption requires an entity wishing to send a secret message to entity A to possess a trusted copy of the public encryption key e_A of A.

Examples of public key cryptosystems include the Rivest-Shamir-Adleman (RSA) encryption scheme [37, 120], and the ElGamal encryption scheme [95].

Padding schemes are often used to prepare plaintext for encryption. Such schemes are used with both symmetric and asymmetric encryption techniques. One example of a scheme designed for use with an asymmetric encryption scheme is known as Optimal Asymmetric Encryption Padding (OAEP). OAEP is commonly used to encode small items, for example, keys. In symmetric encryption systems, padding schemes are often employed when using block ciphers. Padding schemes for block ciphers are discussed in [37].

There are a number of additional differences between the uses of symmetric encryption and asymmetric encryption. These include:

1. The keys lengths for symmetric encryption algorithms are typically relatively short when compared to those for asymmetric encryption algorithms.
2. Symmetric encryption algorithms are typically more efficient at encrypting data than asymmetric schemes.
3. Key management for asymmetric encryption algorithms can be simpler than

that for symmetric schemes, since each entity only needs one key pair, and only the origin and integrity of public keys needs to be assured.

2.3.3 Cryptographic Hash Functions

Cryptographic hash functions have a variety of uses, including as part of most practical digital signature mechanisms. A cryptographic hash function, H, takes a message, m, of arbitrary length as input, and outputs a hash code of fixed length. Generally, a hash function should be easy to calculate.

The basic properties of cryptographic hash functions are [95]:

1. **Pre-image resistance** — for essentially all pre-specified outputs, it is computationally infeasible to find any input which hashes to that output, i.e., to find any pre-image m' such that $H(m') = y$ when given any y for which a corresponding input is not known.
2. **Second pre-image resistance** — it is computationally infeasible to find any second input which has the same output as any specified input, i.e., given m, to find a second pre-image $m' \neq m$ such that $H(m) = H(m')$.
3. **Collision resistance** — it is computationally infeasible to find any two distinct inputs m, m' which hash to the same output, i.e., such that $H(m) = H(m')$.

Examples of cryptographic hash functions include SHA-1, RIPEMD-160 and Whirlpool [37].

2.3.4 Digital Signatures

Digital signatures can be used to provide integrity and authentication services. Let M be the set of messages to be signed and S be a set of elements called signatures. Every entity A that wishes to create digital signatures must be equipped with a private signature transformation $S_A : M \rightarrow S$. S_A should be kept secret by A, and is used to create signatures on messages sent by A. V_A is the corresponding public

verification transformation which maps from the set $M \times S$ to the set $\{true, false\}$. V_A is publicly known and is used to verify signatures created by A. (In practice, S_A is derived from a private signature key, and V_A from the corresponding public verification key). When using digital signatures, the verifier of a signature generated by entity A requires a trusted copy of the public verification transformation V_A of A. Distributing reliable copies of public keys can, for example, be achieved using digitally signed data structures known as public key certificates. A party responsible for generating a public key certificate is referred to as a Certification Authority (CA).

Examples of digital signature schemes include the RSA signature scheme [95] and the Digital Signature Algorithm (DSA) [95].

2.3.5 Time Stamps

A time stamp can be used to record the time of creation or the existence of information [95, p3]. Protocols based on the use of time stamps require synchronised clocks to be available to all parties. In addition secure mechanisms need to be in place to manage the clock synchronisation process.

A time stamp is first obtained from a host clock. This is then cryptographically bound to a message. Upon receiving this message, the receiver then obtains a time stamp from its own host clock. The received message is deemed valid only if it was sent within an acceptance window, and the same message has not been received already. The acceptance window is a pre-determined time interval which makes allowance for the time it takes for the message to arrive, to process the message, and variations in individual clocks.

2.4 Security and Privacy

Providing a privacy service often relies on security services. An example of such a service is confidentiality. When a subject requires that the distribution of data is controlled, i.e. requiring a privacy service, they may use mechanisms that provide confidentiality to prevent unauthorised entities from gaining access to it. This en-

sures that the distribution of data is controlled, and is only available to permitted entities.

Providing a privacy service can also conflict with the provision of other security services. One example is the provision of accountability and privacy. When anonymity techniques are used to protect a user's privacy, there may be no way of holding a user accountable for their actions.

Chapter 3

Privacy and Location-Based Services

Contents

3.1 Introduction **34**
3.2 Threats to LI **36**
3.2.1 Unauthorised LI Use 36
3.2.2 Eavesdropping Threat 37
3.2.3 Unauthorised LI Modification 37
3.3 Using Access Constraints **38**
3.4 Information Reduction **40**
3.5 Using Regulation, Standards, and Guidelines **43**
3.5.1 Regulatory Bodies Providing User Location Privacy 43
3.5.2 Self Regulation 47
3.5.3 Standards and Guidelines 48
3.6 Summary **55**

We begin with an introduction in which we describe of the aims of this chapter. This is followed by Section 3.2 which looks at the threats faced when using LI. The remainder of the chapter is then divided into three sections, each describing the literature associated with one aspect of user privacy in the context of the provision of location-based services. In Section 3.3, we discuss controlling LI by attaching constraints. Section 3.4 discusses controlling the release of LI by changing it before it is sent. Section 3.5 discusses the control of LI through regulation, standards, and guidelines.

3.1 Introduction

This chapter discusses existing research that addresses the user privacy concerns associated with location-based services.

As defined in 2.2.6, privacy enables individuals to have some control over the information about them that may be collected and distributed. This means that they have a degree of control over how this information may be stored, distributed and used.

Privacy Enhancing Technologies (PETs) are tools that are designed to safeguard personal privacy by minimising or eliminating the collection of identifiable data [71]. These tools can protect the identity of a sender or recipient or data. Examples of PETs include blind signatures, allowing authentication without identification [23], anonymous re-mailers, allowing users to anonymously send messages [64], and web-surfing anonymisers [38], allowing users to browse web pages without being tracked. We focus here on means of protecting privacy for LI.

When it is being distributed, information may be changed so the detail that is divulged is at a level that is acceptable to the individual. This can be achieved by reducing the accuracy of the information or by using anonymity mechanisms. Anonymity mechanisms disassociate the individual from the information.

A major privacy concern for LI arises when users do not control its dissemination and use. This is highlighted in a study by Barkhuus and Dey [12]. This study found that users were more concerned about their privacy when a third party was tracking their location as opposed to when they themselves requested a location from a location-aware device. In the latter case, the user is able to control when the location information is gathered and distributed.

The difficulty in controlling information begins when the subject is not in physical possession of it. This can occur in many ways. For example, loss of possession occurs when a user sends LI in return for a service from a third party, or when the user is not responsible for gathering the LI. LI is often gathered by parties other then the user; for example, Enhanced Observed Time Difference (EOTD) technology acquires

LI using base stations, over which the user may not have any control.

Users also seem to have varying perceptions about privacy. Beckwith [14] investigated this using a model proposed by Adams [5]. This model is based on the assumption that the factors affecting user perception of privacy are the identity of the entity receiving information, how the information is used, and the sensitivity of the information. Beckwith investigated user perception of privacy issues arising in a care home using interview techniques, and looked at user concerns about information gathered about them, including LI. The subjects involved in this survey included family members of residents, the residents themselves, and staff members, including management. Beckwith found that subjects had little knowledge about any of the factors discussed in the Adams model. Beckwith concludes that informed consent plays an important role in ubiquitous computing, and suggests the use of easily understandable user profiles would be required to allow users to control information.

Control of information can be achieved in a number of ways. Firstly, a user can simply state how its LI should be managed after it leaves the user's possession. These constraints seem to be the simplest way of achieving control. When a user wishes to transmit LI, he/she simply attaches these constraints to it. However, as discussed earlier, the user is not always the entity which gathers the LI. In cases where LI is gathered by a third party, these user-defined constraints must be made available to it. Additionally, the statements and LI should be accessible in a format which the entities involved are able to understand. Existing work on the use of constraints to control the dissemination and use of LI is discussed in section 3.3.

Secondly, a degree of control may be achieved by reducing the information divulged, and only providing what is necessary for the provision of the service. This is similar to the principle of least privilege, which requires that users are given the minimum privileges to perform a task. This is discussed further, with examples, in section 3.4.

Finally, regulation, standards and guidelines may also be used to encourage entities which hold personal information to handle it in an appropriate manner. Although this does not give the flexibility of control achieved by constraints and information reduction, the user is given a general level of assurance regarding the handling of personal information.

A general discussion about the privacy challenges faced when providing location-based services can be found in [96]. Location privacy concerns in mobile IP networks are highlighted in [85]. The concerns discussed in this latter paper are that users which roam to foreign networks may have LI about them disclosed to other parties. A method of using IPsec [84] in a Mobile IPv6 [35] network to secure traffic between the home network and the foreign network is described in [8]. Secure and private location management for mobile hosts is discussed in [130].

3.2 Threats to LI

In this section, we give an overview of some of the threats arising from the use of LI. The main threats discussed are: unauthorised use of LI by entities that are authorised to possess LI in Section 3.2.1, Eavesdropping Threats in Section 3.2.2, and Unauthorised Modification of LI in Section 3.2.3.

3.2.1 Unauthorised LI Use

One of the main threats to privacy for LI derives from the way that this information will be used. In many cases, the entity in possession of it is permitted to have the LI, but is not permitted to use it in a certain way. Entities that may be permitted to possess LI may include certain service providers, network entities, and entities responsible for gathering the LI.

Two important examples of unauthorised use of LI are the analysis of the behaviour of the LI subject, and the provision of an unwanted service based on this data. Using LI and knowledge of the LI subject, behavioural analysis techniques could be used to analyse the preferences of an LI subject. For example, an employer that is determining the location of an employee during work hours, may also find out their habits outside working hours. When LI has been provided for a service, this information may also be used in the provision of additional services. An example of this arises when LI has been provided to locate the nearest restaurant, and then unrequested advertisements are provided based on this information.

Privacy violations may also take place when an entity in possession of LI passes it on to other entities. Of course, these entities may not have permission to access this LI. Once such an entity is in possession of LI, it again may be used to analyse behaviour and in the provision of unwanted services.

The entity that gathers LI may also pose similar threats.

3.2.2 Eavesdropping Threat

The location of an LI subject may be divulged by eavesdropping on a communications channel. This may occur in a number of scenarios. These include when LI is being transmitted to the LI subject by a third party, when LI is being transmitted to a service provider, or when LI is being transmitted between service providers. Previously we have discussed threats from known entities, i.e. the LI subject itself, the service provider, and the entity that gathers LI. This type of threat differs from the previous Information Leakage threat because the entity receiving the information is not authorised to possess it.

3.2.3 Unauthorised LI Modification

Another potential threat to LI is unauthorised modification. This may take place when LI is first created, when it is being transmitted, or when it is stored.

As discussed in Chapter 1, LI may be gathered by a device possessed by the LI subject or by a third party. LI may be modified in either case. One example of malicious modification of such data by the LI subject is when LI is used as part of an authorisation process. That is, one of the attributes used in an authorisation might involve identifying the location of an LI subject. For example, an LI subject may wish to access a network, and has to be located at a pre-determined location before this is permitted. Modifying the LI in such a case may allow the LI subject to gain unauthorised access to a network.

3.3 Using Access Constraints

We next consider the existing literature on the use of constraints to control the distribution of LI.

P3P [31], also discussed in Chapter 1, is a standardised mechanism for expressing privacy policies. A P3P-enabled browser can be used to fetch the policy of a web site, which can then be compared to privacy settings specified by the user. For example, it can be used to make automated decisions about how to handle cookies. Myles, Friday, and Davies [98] describe a system with an emphasis on usability for user constraints for LI using P3P policies. This system uses a central location server with which requesters interact. The decision regarding whether or not to release LI is made based on the policies. On request, a validating entity, which is previously registered by the user, ensures that the requestor is permitted to receive LI. This entity may also reduce the accuracy of the LI if required by the privacy policy. The resulting LI can then be sent to the requestor. Policies for redistributing this LI are not discussed.

Jiang and Landay [80] discuss the control of information using privacy tags associated with objects. A privacy tag consists of three parts: a space handle, a privacy policy, and a privacy property list.

- The space handle specifies a group to which an object belongs. The boundary of this group may be physical. This can be used to control the movement of a material object. A group with social boundaries can control object movement between different social groups. Activity-based boundaries can control information movement based on an action performed.
- The privacy policy specifies the owner of an object and the entities permitted to perform operations on the object.
- The privacy property list describes the lifetime of the object, its level of accuracy (which may be degraded) when it is transferred to a different space, and the confidence level which applied when the information was gathered.

The two major problems inherent with such schemes are ensuring that the controls described in the privacy tags are adhered to, and that the privacy tags remain attached to the object. Jiang and Landay also discuss the use of a trusted computing base and decentralised servers to store privacy tags to help overcome this problem. The use of trusted computing to achieve privacy is discussed further in Chapter 6.

Snekkenes [132] discusses a language to describe policies. These policies are designed to describe the privacy requirements that apply when LI is requested, or when a service is required by the LI subject and LI must be sent. Expressing controls for LI after it has been distributed is not considered. Additionally, expressing controls on the initial gathering of such information is also not considered.

Duri et al. [42] discuss privacy in the context of automotive telematics. They consider the use of a secure boot procedure, discussed further in Chapter 6, to help ensure the integrity of applications executing within a vehicle's computing environment, as well as using privacy policies to determine whether to distribute LI to entities. Distribution decisions are made by comparing the privacy policy for the subject with that of the service provider. This approach is also taken by Leonhardt and Magee [88], who propose the use of access control policies to determine if access by the querying entity is permitted. Possible inputs to the decision-making process include the time of query, the location of the subject, and the policy of the querying entity.

Spreitzer and Theimer [133] discuss a broad range of security issues related to LI. They observe that access control for LI may be enforced by a location service provider or by a user agent, which may require complex access controls. They also discuss secret groups as a means to handle access control. Users registered to such a secret group may share information freely; however, users that are not members of such groups must abide by access control rules. Authentication mechanisms for entities in the architecture are discussed, providing a basis for access control mechanisms. Additionally, issues with gathering authentic LI are also identified. They conclude that they cannot guarantee privacy without authentic LI, and cannot easily prevent traffic analysis. They also claim that a user agent architecture enables a greater degree of privacy.

3.4 Information Reduction

Approaches to enhancing user privacy by reducing the level of detail in personal information are now discussed. There are two fundamental methods for achieving this. That is, either the accuracy of the LI can be reduced, or anonymity techniques disassociating the individual from the data can be used.

The personal information for which detail is reduced is not always necessarily LI. For example, removing personally identifying information disassociates a subject from LI, which also provides privacy.

Some location-based services, e.g. Tomtom [139] and CitiKey [92], ensure that LI remains private by simply not providing it to third parties. A location may be determined locally, i.e. on a user device, and a service can be provided without transmitting this information. Of course, this can only be achieved when the information required for service provision is all available on the user device. For example, a service may be provided on a user device, e.g. when maps are stored locally on the user device in journey guidance systems. However, this may not be possible where a user wishes to determine traffic conditions on a particular route. The Global Positioning System (GPS), and ultrasonic location systems, are examples of locating technologies which do not require information to be transmitted from a user device. In both cases, a receiver on the user device is used to calculate its location. Further information about these technologies can be found in Chapters 1, 8, and also in [79]. The user device can then use these transmissions to calculate its location. Ultrasonic location systems use numerous fixed transmitters with known locations which transmit signals to a receiver. The receiver then calculates its own location using a combination of the location of the transmitters and the signals received. Examples of systems which work in this way include the cricket system [115], and those discussed in [117] and [69].

Applying obfuscation principles to LI to reduce its accuracy when it is numerical geographical data is easy, at least in principle. This can, for example, be achieved by simply rounding numerically represented data. Reducing the accuracy of this data means that the LI subject will not receive the quality of service that they would had the data been used in its original form. When deciding upon the level

of obfuscation, the LI subject must balance the level of privacy with the quality of service. Of course, techniques that are applied to numerical data differ from those applied to non-numerical data, e.g. where LI is represented as photographic data. Such data may indicate that an individual was at a location at a given time. The privacy of individuals depicted in a photograph may be protected by obscuring the relevant parts of the photo. Indeed, Gibbons et al. [61] discuss blacking out the faces of individuals in the Irisnet sensor network to maintain their privacy.

In some cases LI can be inferred by an entity receiving data, without it having been specifically provided. Privacy issues for such scenarios should also be considered. One example of this type of inferred LI is provided by IP addresses. Transmissions over an IP network require that communicating parties have such addresses. These addresses are often provided on the basis of geographic location. Deductions about the location of a subject may be made from their IP address. The accuracy of such information, however, may not be very high. For example, a subject could use a Virtual Private Network (VPN) to create the illusion that it is at a different location. Even when this is not the case, the IP address may only contain a limited amount of information about the location of the end device. Another example where LI can be inferred is provided by wireless devices; this is because many such devices, for example Radio Frequency Identification (RFID) tags, have a limited transmission range.

An RFID tag is a small, cheap, wireless transmitter which can be used to transmit small amounts of information such as a serial number. This serial number can be used to uniquely identify a product, and such devices have numerous possible applications. For example, in a supermarket, RFID tags may be used instead of bar-codes to identify products, e.g. for stock-taking or at the point of sale. Other suggested applications include fridges which identify when a product has expired, and washing machines which identify the suggested wash cycle for items of clothing. These tags also pose a threat to privacy. Thieves could, for example, use the information provided by the tags to decide whether or not to steal a product stored in a car. More generally, a passer-by can identify information associated with any product containing an RFID tag. Various approaches have been suggested to reduce these risks. Tags may be rendered totally inactive by disabling them; however, this also eliminates potentially useful applications. Methods for preventing signals from

reaching third parties have been proposed, for example jamming signals or using a shield through which radio waves may not pass; however, these techniques have similar drawbacks to the previous approach. Juels, Rivest and Szydlo [82] propose the use of a blocker tag that simulates a subset of the total set of serial numbers by transmitting this subset. The RFID reader is unable to distinguish the serial number of the actual RFID tag from the ones being transmitted by the blocker tag. Using this approach, Juels, Rivest and Szydlo discuss privacy policies which may be used to obscure only certain serial numbers. These may be serial numbers which are attached to products, the possession of which a consumer does not wish to advertise.

Proxies may also be used to enhance user privacy. Escudero-Pascual and Maguire Jr [49] suggest using a proxy to make requests to an LBS provider for a number of different locations. This prevents the LBS provider from determining the exact location of the requestor. Zhu, Mutka and Ni [154] also suggest the use of proxies; however, the focus of their work is to provide authentication, confidentiality, integrity and non-repudiation between entities, and they do not discuss policies for access control to LI.

Beresford and Stajano [17, 18] discuss the use of mix zones to provide privacy. The concept of a mix zone is derived from a mix network [24]. A mix network uses a number of mix nodes to provide unlinkability of correspondents to an observer. The routing information is layered, showing only information about the next node. The level of anonymity is defined by an anonymity set, which is, in this case, the number of distinct senders and receivers. A mix zone divides an area into zones, preventing observers from tracking the locations and movements of users. This is achieved by mixing the identities of users using a proxy. Beresford and Stajano also suggest a method for quantifying levels of anonymity using the number of people found in the mix-zone and the probability that a subject entering a particular zone will exit towards a certain other zone.

Even when anonymity mechanisms are used, it may still be possible to estimate the movements of subjects. The use of Bayesian Filtering techniques to achieve this is discussed by Fox et al. [53].

Gruteser and Grunwald [63] discuss providing anonymity by reducing the accuracy

of LI, as well as the time at which it was gathered, and providing values from a set instead of precise information. The level of anonymity this approach provides is discussed in terms of the minimum number of entities within a set of locations and times. Once the defined minimum number of entries exist within the anonymity set, the level of accuracy reduction in time and location data is provided to the requester of LI. The time, calculated as a random time within the anonymity set, and the LI, indicating a random location within the anonymity set, is also divulged. The authors acknowledge that this method is not always practical, particulary when LI is required quickly, as the requester must first wait for the anonymity set to be filled with the minimum number of entities.

3.5 Using Regulation, Standards, and Guidelines

This section is divided into three parts. Firstly, we discuss regulatory bodies which deal with LI. In particular, we discuss the legal aspects of dealing with LI in a variety of jurisdictions.

Secondly, we consider self-regulation. This has been proposed as a means of encouraging the uptake of services based on LI. In this approach a corporation provides assurances about its code of practice for the handling of LI, with the goal of increasing end user confidence.

Thirdly, we discuss various standards and guidelines targeting user privacy for LI.

3.5.1 Regulatory Bodies Providing User Location Privacy

The regulatory aspects of location privacy vary around the world. Regulations also play a significant role in the provision of a variety of services. Ackerman, Kempf and Miki [4] argue that regulation has played an important role in the successful deployment of LBSs in Japan.

Laws regulating LI differ widely for historical and cultural reasons. We now discuss regulations relevant to the privacy of LI for some of the most significant users: the

United States (US), the European Union (EU), and Japan.

3.5.1.1 US Regulation

User privacy for LI in the US is covered by a number of different laws. The result is that there is no single document describing how privacy issues for LI should be handled. This section looks at the history of, and the issues raised by, the relevant laws, and examines their current standing.

Customer Proprietary Network Information (CPNI), as discussed in the Telecommunications Act of 1996 [1], is information gathered about a customer by telecommunications companies, such as telephone and mobile phone operators. This information includes details about phone calls made by customers, such as the phone number called and the duration of the call. The Act includes provisions that protect this data by requiring prior customer consent before CPNI can be used. The aim of this requirement was to prevent telecommunications companies using this data for marketing purposes and selling it to other entities without customer consent. In 1998 the FCC stated that, when implementing this law, prior consent means that customer consent must be gained through written, oral, or electronic means [4]. This was then challenged in 1999 by a company called US West, together with other concerned parties, in US West versus Federal Communications Commission (FCC) [4]. US West argued that an opt-in rule, where a telecommunications company is required to obtain approval from customers before using CPNI, was a violation of the First and Fifth Amendments of the United States Constitution. The court concluded in 2000 that the FCC regulations were impermissible. However, they did not order the FCC to adopt an opt-out approach, and instead discussed less burdensome alternatives.

The Cellular Telecommunications and Internet Association (CTIA) represents all players in the wireless community. The CTIA believed that LI was too significant to be a part of CPNI. They proposed that the use of LI should have separate rules to CPNI. In particular, the CTIA wanted customers to have prior notice, and an opportunity to provide consent, before LI was collected or used, assurances about the security of collected data, and for the rules to be technology-neutral [4]. These proposals were rejected by the FCC. The FCC eventually gave an ambiguous ruling

that opt-in approval was required when information was provided to third parties which did not provide communications-related services. Finally, note that the Telecommunications Act of 1996 does not address recent developments in technologies where standard phone numbers are not used, such as the Voice Over Internet Protocol (VOIP), RFID tags and Wireless Local Area Networks (WLANs).

The Wireless Communications and Public Safety Act of 1999 amended the Telecommunications Act of 1996, stating that LI should be provided for emergency calls. However, when used for a purpose other than an emergency, prior consent was required, i.e. users must opt-in. Although this contradicted the US West versus FCC case, the FCC stated that this opt-in scenario only applied to this part of the rule. The FCC also mandated that opt-in was required when CPNI is sent to a third party and where a service was not provided [4].

In an attempt to set rules covering LI, Senator Edwards proposed legislation restricting the ability of a location-based services provider to collect, use, retain, or distribute location data for their customers without prior notification and consent [44]. This proposed legislation stated that the opt-in rule should be used for LI, with the exception of data requested by court order, aggregate data, and data used for emergency services. However, this was not passed into law.

In addition to rules set by the FCC, individual states also have their own rules about LI. For example, in Washington, the Utilities and Transportation Commission requires an opt-in rule when LI is shared with a company that is not commonly owned [4].

Another US law affecting LI is the Digital Millennium Copyright Act (DMCA) of 1998. This law requires that consumers are locatable so that they can be prosecuted when violating this law. This would mean that concealing location information is illegal. This rule also varies depending on the state in which you are located. For example, in Arkansas, it is only illegal to conceal LI during a criminal act. In South Carolina, it is only illegal to conceal LI if there is proven fraudulent intent.

3.5.1.2 EU Regulation

The Organisation for Economic Cooperation and Development (OECD) issued guidelines for privacy protection in 1980, which stated that data subjects have the right to correct their personal files. Additionally it also required that any action based on this personal information which has not been previously agreed required further permission [4].

Members of the EU have guidelines for privacy set by the 1995 Data Protection Directive, and they then model their own laws around this. In addition to protecting privacy, transfer of private information to non-EU countries which do not provide suitable privacy protection is also covered by this directive. Generally, the directive states that retention and use of data not required by contract or legal reasons requires consent. Also the directive states that subjects should be able to have access to their information and correct errors, and be able to find out from where the data originated. They should also be permitted to opt-out and have the right to take action when it is used unlawfully. The 2002 Directive on Privacy and Electronic Communications provides a technology neutral set of guidelines also covering LI. It states that LI used for services other than those for which it has been previously agreed requires further permission. It also states that it should be possible to revoke a permission after it has been given. When requesting permission, details about the intended use should be provided. When data is kept for billing purposes, it should be deleted or made anonymous after its use.

The directive provides guidelines for the member countries, and it is up to the individual countries to then implement these as law. The directive has been interpreted differently by different countries, and some countries have been slow to implement the necessary laws, an important factor in the use of location-based services [4].

3.5.1.3 Japanese Law

In Japan, the 1988 Act for Protection of Computer Processed Personal Data held by Administrative Organisations protects against misuse of data collected by governments, which includes LI. It does not cover data which may be used for the

purposes of national security and law enforcement. This act is based on the OECD guidelines discussed in section 3.5.1.2. The act allows data subjects to view their data, find out who controls it, and the purpose and method of its collection.

For other organisations, the 2003 Personal Data Protection law states that subjects should be informed about data collection, although this law does not specifically discuss LI. Subjects must be notified of the intended use of data, and consent must be acquired before data is sent to third parties. Moreover, data subjects should be able to view, correct and withdraw consent, reasonable measures should be taken to secure data, and there should be the opportunity to complain when personal information is misused.

The Ministry of Posts and Telecommunications has also published guidelines on the protection of personal data in the telecommunications industry. These guidelines are not binding. The guidelines cover the protection of personal data during transmission, and may be used as the basis for further regulation. They state that LI should not be disclosed unless the subject has given prior consent. Exceptions to this include when presented with a warrant issued by judge, a life threatening emergency or other legal grounds.

3.5.2 Self Regulation

Service providers may maintain a privacy policy for personal information. This provides information on how providers intend to handle such information. Indeed, service providers which handle LI may also have specific policies for this type of personal information.

One example of such a policy is provided by Cingular Wireless [28], a wireless services provider based in the USA. The Cingular Wireless Privacy Policy is reproduced in appendix A.2. Cingular offer a variety of optional services in their 'mMode' package, including Internet access using the Wireless Access Protocol (WAP). Most importantly, Cingular's mMode service can provide personalised information based on location. Further information about mMode can be found in [27], and this information is reproduced in appendix A.3.

3.5.3 Standards and Guidelines

We now look at various standards and guidelines covering the privacy of LI resulting from the Geographic Location/Privacy (Geopriv) charter [73], maintained by the Internet Engineering Task Force (IETF) Secretariat [6]. The IETF Geopriv working group is responsible for producing documents based on the charter. It aims to produce standards to enhance security for LI in IP-based networks. More specifically, the working group investigates authorisation for the use and dissemination of LI, together with integrity and privacy requirements. It also looks at requirements for representing this data.

The IETF produces documents known as Requests For Comments (RFCs) and Internet-Drafts. Internet-Drafts represent work in progress, and are a means to share ideas with interested parties. Internet-Drafts are usually valid for a maximum of six months. RFCs are documents which must first be published as Internet-Drafts. They form a set of technical and organisational notes about the Internet. RFCs produced under the Geopriv charter to date are [33, 34, 108, 109, 112]. Current Geopriv Internet-Drafts are [93, 125, 126, 127, 128, 138, 143, 150].

The following six subsections cover the output of the Geopriv working group. The Geopriv requirements describe the general requirements for the Geopriv protocol, including those for security; these are discussed in section 3.5.3.1. A presence architecture, describing the availability of an entity, is discussed in section 3.5.3.2. DHCP in relation to LI is described in 3.5.3.3. A document format for describing a Location Object (LO) is described in section 3.5.3.4. An algorithm for resolving conflicting rules is described in 3.5.3.5. Finally, the application of the RADIUS protocol to Geopriv is described in section 3.5.3.6.

3.5.3.1 Geopriv Requirements

General requirements [33] for the Geopriv protocol cover the means to securely gather and transfer LI within the scope of a Location Object (LO). An LO is a data structure used to distribute LI. To facilitate LI privacy, requirements for secure transmission of an LO, user controls for an LO, methods to reduce the accuracy

of LI, a core set of rules carried within the scope of the LO, and enabling user anonymity are all considered.

The main entities within the Geopriv protocol are: the Location Gatherer (LG), responsible for determining LI and creating the LO; a Rule Holder (RH), responsible for maintaining rules for receiving, filtering and distributing LOs; a Location Server (LS), an entity which receives LOs from the LG, applies rules from the RH, and then distributes the LO according to the rules; and, finally, the Location Recipient (LR), which is the end recipient of an LO.

The Geopriv protocol uses a set of privacy rules to regulate the collection, use, disclosure and retention of LI. Additionally, these privacy rules may be used to describe any changes made to LI prior to its distribution. For example, the accuracy of the LI may be reduced. The scenarios when these changes are applied are also described in the privacy rules.

Anonymity, unlinkability and the use of pseudonyms are also noted as requirements for privacy. Geopriv regards anonymity as a means to ensure that a user is unidentifiable; unlinkability is defined in [33] as a means of ensuring that colluding entities cannot learn that multiple resources have been used by the same user, and pseudonyms are unique identifiers. Pseudonyms may additionally be unlinkable to a user. This allows an entity to link the activities of a pseudonym; however, the entity still cannot identify the actual user. Geopriv also considers message authentication and end point entity authentication. Unfortunately authentication mechanisms are not usually compatible with privacy mechanisms, and may require that the identity of an entity is easily linkable.

Additional security considerations arise when the LO is transferred. Also discussed are structural requirements for the LO.

Threats to the Geopriv protocol are identified by Danley et al. in [34]. The main threats identified are eavesdropping in a network and impersonation, where a malicious party may pretend to be a valid entity in a network. To address these threats, confidentiality between entities, and authentication and authorisation between entities, are identified as security requirements. Other threats discussed are denial

of service attacks, attacks on stored data, and policy information as a privacy violation. Limiting the distribution of policy information is identified as a means of overcoming this hazard. Additionally, short-lived identifiers and pseudonyms also provide some additional privacy for users. Denying frequent location requests prevents unnecessary gathering of tracking information. The use of default policies is also identified as a means to reduce threats.

3.5.3.2 Presence Architecture for LO distribution

The Presence Information Data Format (PIDF) [136] was developed as a means to communicate private information. This format is extended in [108], [109] and [138] for its application to Geopriv. Presence allows one entity to determine the availability of another for communication. The extensions needed to the PIDF to accommodate Geopriv are discussed in [138]. This latter document describes the PIDF within the scope of an XML schema to form an LO. In this case, the LO is described as a Presence Information Data Format Location Object (PIDF-LO). The scheme suggests the use of the Geography Markup Language (GML) [77] to describe geographic data. This data will not remain static when a subject is moving. A filtering method to determine if these changes are significant enough to warrant notification is discussed in [93].

Three fields in the PIDF-LO can be used to limit its retransmission and retention, together with a reference to external rule-sets.

Retransmission of a PIDF-LO may either be permitted or denied. Finer granularity in this retransmission control process would allow a more flexible model. Other limitations on retransmission may be imposed based on the type of service required by the end consumer. Additionally a 'white list' containing entities to which LO may be transmitted may allow the end user further flexibility. Such a white list may be transmitted within the scope of the PIDF-LO, or held in a server.

The field describing retention limits contains information about the length of time for which the PIDF-LO may be retained. This can be used in conjunction with the value of the time stamp field in the PIDF-LO to determine how long the LO can

be retained. If the time-stamp field is empty, the LR can use the time at which it receives the PIDF-LO in its place. A greater degree of privacy can be provided if an expiry point is stated rather than a time stamp and a time interval. This would mean that LRs could not establish the precise location of an end consumer. Potential issues here are that the use of different time-zones may cause confusion. A solution for this may be to use the LI to establish the time-zone of the consumer and to use this. Alternatively, the time-zone of the originator may be described in an additional field. Another alternative is to use a global time standard such as Coordinated Universal Time (UTC) [78] for specifying time.

The rule-set reference field may describe the location of the rule-set. Access to this may require use of an authorisation and authentication procedure.

Mechanisms for constraining and interpreting data in the PIDF-LO are discussed in [150]. The use of standardised data formats for LI may reduce ambiguity when representing LI. Mechanisms to describe the accuracy of LI are also described. One of the representations discussed is 'arc band', which uses Timing Advance (TA) information. This information is used to define two concentric arcs, and the precise location is between the two arcs. Alternatively, a polygon can be defined, i.e. the precise location is within a set of points forming a shape.

3.5.3.3 Dynamic Host Configuration Protocol (DHCP) for Geopriv

The Dynamic Host Configuration Protocol (DHCP) [40] is used by clients in a TCP/IP network [135] to obtain network configuration information, facilitating communications with other devices in the network. Examples of configuration information include the unique network IP address, a subnet mask, the unique IP address for a DNS server, and information about the default gateway. Further examples can be found in [40].

Within the context of Geopriv, DHCP can be used to provide coordinate based LI using a DHCP server [112]. One of the major motivations for locating an IP address is to locate the origin of an emergency call from an IP phone.

The information provided could then be translated into LI, which may then be transmitted to the client. The LI is represented as a pair of latitude and longitude coordinates, for which various standards are discussed. Unfortunately, security is not considered in DHCP, so the use of this protocol to transmit LI poses potential security risks. To overcome this, proposals have been made [125] to reduce these risks. One possible approach is to transmit LI only when requested, in order to limit its distribution.

Also discussed in [112] is the type of LI to be used, and the nature of the device to be located. For example, this may be the actual DHCP client or the DHCP server.

3.5.3.4 A Document Format for Expressing Privacy Preferences

Issues surrounding the expression of privacy preferences are discussed in [126] and [128]. The focus of [126] is the issues which arise when the privacy rules given in [128] are combined with those established for presence, specified in [121].

After an entity requesting LI has been authenticated, a permission-combining algorithm is used to permit or deny access to LI. This algorithm uses the common policy to derive a permission by evaluating its rules. If access to LI is permitted, the rules may also state that changes must be made to the LO, which must be applied before the LO is sent. The common policy rules are expressed in the form of an XML schema. The rules described can only provide permissions (i.e., they are not permitted to express denials of access rights). Such a requirement means that rules do not have to be ordered, which overcomes the problem of conflicting permissions. The permission granted to the requester is the intersection of all the rules. The rules are made up of three parts, i.e. conditions, actions and transformations.

A condition is a statement which may evaluate to either true or false. Such a statement may ask questions about the identity of the requester or the validity of the LI. Actions and transformations specify how a request is handled after the requirements of the conditions have been evaluated. For example, an action statement may state that if LI is sent to an entity, then the subject of the LI must be notified. A transformation describes how LI data must be modified before transmission to the

requester.

The identities of the requester and the owner of the LI are used for various authentication procedures. Generally, before transmitting the LO to a requester, the rule-set must be removed, unless otherwise stated. The concern here is that the rule-set may disclose information about entities trusted by the owner and requester.

The minimum security mechanisms provided by Geopriv include mutual end point authentication, data object integrity, data object confidentiality, and replay protection. Secure/Multipurpose Internet Mail Extension (S/MIME) [43] provides a means to securely transmit MIME [54] data. This provides an extensible method of transmitting data encoded in a variety of character-sets. As specified in [110], S/MIME is to be used to transmit LOs [128], and should also ensure the protection of the rule-set. Ensuring protection against eavesdroppers is also a requirement. The protocol suggested to achieve this is https. Authorisations may be made based on the geospatial location condition.

3.5.3.5 Rule Evaluating Algorithm for Combining Permissions

In some cases, there may be more than one rule in a rule set describing the same attribute. For example, one rule may state that access to LI may occur within a given time period; however, another rule may conflict with this. To resolve this potential problem, a simple algorithm [126] was devised to overcome conflicts without reducing the level of privacy.

First, if there is a rule within the rule set in which the data-type is boolean and equates to true, then the combined value of this rule is said to be true. When the data-type is an integer, the highest value is said to be the combined value. Data which is undefined is not included. Finally, when the data-type of the rule is a set, the union of all the data is said to be the combined value. Undefined data is not included in this process.

3.5.3.6 Carrying LOs in RADIUS

Authentication, Authorisation, and Accounting (AAA) protocols provide each of these three services. Authentication, Authorisation, and Accounting are defined in Chapter 2. Remote Authentication Dial In User Service (RADIUS) [119] is an AAA protocol, developed by the IETF, that is designed for use between a client and a RADIUS server. There are a number of applications for RADIUS, including IP mobility and network access. Information gathered at a RADIUS server can be used for a number of purposes, such as billing and the generation of usage statistics.

The distribution of Location Information in RADIUS is considered in [143]. This provides a means to facilitate location aware billing. This can be used to provide network access and charge a client appropriately, taking into account factors such as tax. This is achieved by adding the following attributes to the RADIUS protocol: Operator-Name and LI. The Operator-name attribute uniquely assigns a name to a network based on the operator type. Operator-Namespace, which is a part of the Operator-name attribute, identifies the owner of access network. This may use identifiers belonging to mobile networks.

The LI may be either civic, which identifies information such as the street or city, or geospatial, which identifies latitude, longitude and altitude. The LI attribute in this case also determines whether the LI is for the AAA server or the end user.

A user may hide their privacy policy information and release it only when a requesting entity is authenticated. Alternatively, a user may simply have a default policy which may be distributed freely. Such policies may also indicate the type of LI that the device is capable of transmitting. This may be civic or geospatial.

The visited network is not permitted to distribute LI unless otherwise stated in the privacy rules. Networks must follow the basic rules of the home network if these exist. If there are extended policy rules, then these should be available from the home network. These policy rules must be retrieved from the home network and obeyed. If the RADIUS client learns the basic rules independently, then these rules must be sent when access is requested. The rules must be evaluated before LI is sent. This also applies to extended rules. If an entity is in its home network then, given

that the trust relationship between the entity and the home network is potentially strong, this network can be used to store policies. By default, these policies must not be distributed to third parties. Basic policies must be distributed to the visited network. This is also the rule for extended policies. Additionally, legal obligations must also be met.

3.6 Summary

This chapter has described the threats to privacy that apply to LI. These threats will be used in the following chapters to motivate the discussions.

We have also provided an overview of current technologies designed to address some of the threats described. The regulations in place to protect LI in Japan, Europe, and the USA, the three biggest consumers of location-based services, were also discussed.

The most significant piece of work in this area is that of the Geopriv working group. One concern with this work is that its focus is on Internet applications, and that it may not apply to other network architectures. Limitations with other existing schemes that were discussed motivate the chapters that follow.

Chapter 4

Constraints and Location Based Services

Contents

4.1 Introduction **57**
4.2 A Model for the Use of LI **58**
4.2.1 The Roles 59
4.3 Using Constraints with LI **61**
4.3.1 Constraint Types 62
4.3.2 Uses of Constraints 64
4.4 Limitations of Constraints **66**
4.4.1 Difficulties in Preventing and Detecting Constraint Abuse . 66
4.4.2 LI Constraint Predicaments 67
4.5 Combining Constraints with Auditability **68**
4.6 Specification of Constraints **70**
4.7 Summary **71**

This chapter considers the possible use of constraints to control the dissemination and use of location information (LI) within a location-based service architecture. The various types of constraint which may be required are also considered. Finally, issues and risks with the possible use of constraints are discussed, as are possible solutions to these hazards. Much of the work described in this chapter has previously been published in [57, 59].

This chapter begins in section 4.1 by introducing the concept of constraints in the context of location-based services. Section 4.2 introduces a model within which LI may be used. This model helps us to better understand the requirements for the use of constraints in the context of location-based services. We then go on in section 4.3 to describe the types of constraints that may be used. Possible limitations of the use of

constraints are described in section 4.4. The use of auditing functions is introduced in section 4.5. This leads to a possible means to prevent the unauthorised distribution of LI through co-operation amongst users of LI. In section 4.6 we also discuss the advantages of a standardised language for describing constraints and LI, and briefly look at a possible candidate for such a language. The chapter is summarised in section 4.7, which also contains a review of further work which may aid the wider use of location-based services.

4.1 Introduction

This chapter introduces the notion of LI constraints; these constraints are a means by which an LI subject can exert control over the distribution and use of its LI. Existing techniques for expressing constraints for Location Information include schemes proposed by Schulzrinne et al. [126] and Myles, Friday, and Davies [98]. Schulzrinne et al. look at means of expressing privacy preferences for transmitting LI in the Geopriv protocol. They consider placing controls on the retention of LI, entities to which it is distributed, and how the accuracy can be changed before distribution. They do not consider imposing constraints on how LI will be used. Myles, Friday, and Davies consider a centralised server that holds a user's LI privacy preferences. While their proposed scheme provide a mechanism for access control for LI, it does not consider a receiver that wants to redistribute the LI that is retrieved. Securely attaching the constraints to the LI, together with distribution constraints, as described in this chapter, provides a means of achieving this. Neither of these existing proposals consider entities other than the user that might benefit from use of LI, e.g. an authorisation mechanism that requires a user to be located at a certain place before access to a resource is granted.

In the context of this chapter, LI constraints are simply rules associated with a specific piece or set of LI, restricting the ways in which the associated LI may be used and/or disseminated. To be effective, the constraints must be bound to the associated LI, typically by cryptographic means when the LI is in transit, and by access control techniques for stored LI.

LI constraints can be used to help manage the use and distribution of LI. We

investigate the possible constraint requirements which an LI subject may have, and discuss how these may be fulfilled. By looking at various uses for LI we investigate restrictions which may be placed on these uses. We look at the limitations which can be placed on the distribution of LI, and how responsibility might be determined when LI constraints are abused. Constraints may also be placed on the storage of LI, whereby an LI subject may be able to limit the amount of time for which an entity can hold LI.

The limitations of using constraints to control LI are also studied. Although constraints may allow an end user to have some degree of control over its LI, placing constraints on LI also allows an entity to gain additional knowledge about the LI subject. For example, in order to place storage constraints on LI, time stamps, as discussed in Chapter 2, may be used to limit the length of time that an entity is permitted to store the associated LI. However, this mechanism also poses a risk, since placing a time stamp on LI might allow receiving entities to learn the time that the LI subject was at a particular location. We discuss an alternative mechanism which involves recording the time at which LI expires, instead of using a timestamp and duration.

Similarly, statements about where LI must not be distributed passes information to the receiving entity about entities which the LI subject may have an aversion to. Contrariwise, statements about where LI can be distributed and ways in which it can be used may also give information about services which an LI subject uses. These issues are revisited in section 4.4.2.

4.2 A Model for the Use of LI

We start by defining the entities involved in a location-based service architecture. The relationships between the various entities are also described. We then briefly consider the issue of abuse of LI.

4.2.1 The Roles

- **Malicious Party.** This is an entity with malicious intent. A malicious party may act as a threat to the confidentiality, integrity or availability of LI for one or more LI subjects.

- **User Device (UD).** This entity is a device with which the LI subject may interact, e.g. to invoke a location-based service. Such a device may either be static, e.g. a desk-top computer, or more typically mobile, such as a mobile phone or Personal Digital Assistant (PDA). It is, in fact, this device regarding which LI is generated rather than the user him/herself, since there is typically no way to directly measure the location of individuals. Thus this entity is a key part of the model.

- **LI gatherer.** This is an entity which gathers or possesses LI about an LI subject. A GPS receiver is an example of an LI gatherer, as it obtains location data. An entity in a GSM network which keeps signalling data for a UD is also an example of a LI gatherer. Although a GSM network does not normally pass on this LI (except in certain special cases), it certainly possesses such information, and could, in an appropriate environment, be a valuable source of LI for commercial use.

- **Location-Based Service (LBS) Provider.** This entity provides a service based on LI. This could, for example, be a multimedia location-based service, e.g. a vehicular navigation, gaming or advertising service.

- **LBS directory.** This entity provides information regarding the LBS providers which are available for use by a particular user. The LBS directory may itself use LI regarding the service consumer when providing the service. For example, it may show a service requestor lists of LBS providers providing information about particular types of retail premises in the area of the requester.

- **Network Entity.** This is a component which provides a network service to a UD. Two important types of Network Entity are the local base station which provides network access to the UD, and the UD's 'home network' with whom the UD owner has a contract and charging arrangement for the provision of network services.

- **Regulator/Legal authority.** This is an entity which exerts legal or regulatory control over the management and use of LI. This includes telecommunications regulators, data privacy authorities, law enforcement bodies, and auditors.

In the roles above, the LI subject must trust the LI gatherer to distribute LI according to their wishes. This is not a problem when the LI gatherer is a GPS receiver and the UD is a mobile phone, because the LI subject has physical access to this information and can control how it is distributed. However, when a network entity is responsible for this task, as is the case with EOTD and location systems based on 802.11 wireless networks, the LI subject has less control over its LI. Network entities may be regulated, in which case private data must be controlled in a certain way. See Chapter 3 for further information. Although regulations set by a legal authority must take precedence, the LI subject may want to set additional controls on the way that their LI is handled.

Although certain LBS providers may be permitted to have LI, the LI subject may want to control the way that they use and redistribute this information. Additionally, they may want to restrict how long they keep this information. Also, although the LI subject may trust the LBS provider that first receives the LI, they may not trust the entities to which the LI might then be redistributed. Even when the LI subject trusts the entity to which it first sends LI, it may still want to set restrictions on how its LI is used, stored and redistributed. For example, an LI subject that requests a navigation service from an LBS directory may want to describe ways of reducing the accuracy of its LI when it is sent to certain other navigation services. This might, for example, be because the LI subject's trust relationships with navigation service providers vary.

Example Scenario

Suppose Alice is employed by XYZ technologies as an IT support engineer. Alice spends most of her day going to user PCs to resolve technical issues. As part of the company's drive for efficiency, they ask that Alice can be located during the day by her team. Her team can then use this information to send her to the nearest user

that is having a technical problem. She owns a GPS-enabled mobile phone. That is, her UD also contains an LI gatherer. Her company's access control requirements state that she can only login to her computer when she is physically located inside the building. Alice does not want to be located by people other than her team, and does not want her LI to be used for any purpose other than logging on to her computer. Additionally, Alice's ex-husband, Marvin, who also works with XYZ technologies, should not be able to locate her.

Once Alice leaves work at 17:30, she does not want to be located by her company or Marvin. When she leaves work she decides that she wants to go to the cinema. She does not know the area very well so she uses the web-interface on her UD to access an LBS directory. She provides LI and then uses this to find cinemas that are closest to her. She does not want the LBS directory service to know exactly where she is, only that she is in a certain area.

On her way to the cinema, she sees a car accident. She quickly calls the emergency services; however, she does not know the name of the street she is located in. As required by regulation for the country she is in, the network entity that interacts with her mobile phone calculates her location and passes it on to the emergency services.

4.3 Using Constraints with LI

As discussed in Chapter 3, the major threats to LI come from its unauthorised use, unauthorised access through eavesdropping, and unauthorised modification. In this chapter, we aim to describe ways of reducing unauthorised use. By placing certain controls on LI, we can ensure that it is only used in a way that is acceptable to the user. We refer to these controls as LI constraints. These controls are designed to limit the use, distribution and storage of LI. Observe that controlling all the entities that receive LI, and the duration that LI is stored by an entity, will also control its use. An LI subject may not want its LI to be used by entities that it considers untrustworthy. The use of LI distribution constraints may help to achieve this. Additionally, the LI subject may not mind an entity using its LI for a day, but may not want its LI used by that entity after that time.

We refer to any violation of LI constraints as LI constraint abuse. LI constraint abuse is *any use, distribution or storage of LI which contradicts the rules defined by the constraints.* We also assume that the LI, together with the constraints, is transmitted and stored in a secure manner. By this we mean that the receiving entity has been authenticated by the sender, that LI confidentiality and integrity are not compromised during transmission, and that the LI constraints are securely bound to the LI. This is discussed further in Chapter 5. In order to set constraints on LI, we must first look at how an LI subject may want to restrict the use and distribution of LI.

Depending on the access rights of the LI receiver, the accuracy of LI may be degraded. For example, if a potential receiver does not meet all the rules set in the constraints, then it might be sent a degraded version of the LI. This degraded LI could provide significantly less accurate information about the location of the LI subject than is available in the non-degraded LI. For example, it may only provide the town in which the LI subject is located, instead of the street.

Also, constraints should be in some common format which is automatically processable. In this section we discuss how constraints and LI can be used together.

4.3.1 Constraint Types

The types of LI constraint which may be required are now considered.

4.3.1.1 Storage Time Constraints

Storage time constraints may be used to limit the duration that an entity can store a particular user's LI. This can be done in two ways.

The first method relies on the use of time stamps. A time stamp, as discussed in Chapter 2, can be used to record the time of creation of LI. By adding a validity period, a statement can be made that an entity should not hold LI subsequent to its expiry. For example, the LI subject may state that an entity cannot use LI

beyond one hour after the time set by the time stamp has elapsed. The use of time stamps requires additional security mechanisms. Time stamp based protocols require synchronisation and secure clocks. Also, there should be secure mechanisms for the relevant entities to synchronise their clocks. Synchronisation and secure clocks are discussed further in 2.3.5. The fields needed for such a constraint would be the issue date/time and the validity period.

The second method of adding time constraints to LI is by stating the time at which LI expires. This could be in the form of a date and time after which LI cannot be held. This would eliminate the need for a time stamp. The entity receiving LI will, however, need access to a clock which has been synchronised with the entity generating the constraints, in order to learn when LI is invalid. The field necessary for this scheme would be the expiry date/time.

4.3.1.2 Distribution Constraints

An LI subject may want to constrain the distribution of LI. Distribution constraints can be specified inclusively or exclusively. Inclusive constraints would show the entities who are permitted to possess LI. Exclusive constraints would show the entities who are not permitted to possess LI.

Consideration should also be given to the way in which LI distribution is managed, and which entity is accountable for misuse of LI. This could be the entity which sends LI to an entity who is not permitted to receive it, or it could be the entity which receives and then stores LI when it is not permitted. Of course having both sender and receiver responsible for protecting LI would be most desirable, to ensure that the probability of misuse is reduced, or at least that the probability of misuse detection is maximised.

4.3.1.3 Usage Constraints

An LI subject may want to place constraints on the use of LI, to allow it to restrict the way in which their LI is used. Difficulties when constraining usage arise when

attempting to enumerate all the different applications of LI, because of the wide range of possible uses. An attempt to classify the main possible uses of LI is given in section 4.3.2 immediately below.

4.3.2 Uses of Constraints

The uses of LI can be divided into two main types. LI can be used to:

- provide the LI subject with a service or with location details, or
- provide a service or location details to a separate entity.

The LI subject may, of course, not wish other entities to gain access to its LI, and hence may use constraints to limit uses of LI falling into the second category. Both these two main categories can be further sub-divided. Some examples of uses of LI within these two broad classes are shown below — see also Figure 4.1. These lists are not meant to be exhaustive.

4.3.2.1 Providing Services to the LI Subject

We consider five main categories for the use of LI for which the benefiting entity is the LI subject.

- *Location-based security* may be used to provide a security service to the LI subject. For example, the LI subject may not want to carry out transactions with retailers from certain locations. The LI subject can check the location of the retailer and use this information to decide whether or not it wants to carry out a transaction.
- *Location-based messaging* may be used to inform an LI subject regarding any "buddies" who may be located nearby. These buddies may be a members of a list of people maintained by the LI subject. This is similar to Internet messaging with the added location property. An LI subject may not want

"buddies" to know his/her location at a given time, a restriction which can be supported by the use of appropriate LI constraints.

- *Navigation* services are currently one of the most commonly proposed uses of LI. LI is used to locate the LI subject, and information is provided according to the subject's navigational needs. For example, the navigation service may have information about nearby traffic congestion which should be avoided. It may also plan a route which avoids this congestion.
- *Directory services* may be used by the LI subject to find local services. For example, LI may be used to help the LI subject locate nearby restaurants.
- *Other services* may be provided to the LI subject, including weather services, where weather information is provided based on the LI subject's location.

4.3.2.2 Providing Services to Other Entities

We next consider four categories of use of LI for which the benefiting entity is not the LI subject.

- *Advertising* based on location [148] is a potentially useful tool for retailers. For example, as an LI subject is passing a shop, messages about special offers may appear on the UD. Of course, advertisements may not be something that an LI subject wants; this is a problem similar to junk mail. LI constraints could be employed by the LI subject to prevent such a use being made of LI.
- *Location-based security,* as its name suggests, refers to the provision of a security service based on the location of the LI subject. In particular, access control and authentication [36] may be provided based on the location of an LI subject. An example of where this may be useful arises in the context of wireless LANs. So called 'war-driving' attacks allow unauthorised users to access a network from outside the perimeter of a building [21]. A secure mechanism for providing LI to the entity controlling the network would prevent such attacks. An example of such a mechanism is discussed in [124].
- *Location-based safety* describes the situation where, in an emergency, details of an LI subject's location are sent to the safety entity so that the subject can

be located quickly and efficiently. Although the LI subject may eventually be the beneficiary of this service, its immediate use is by the emergency service. In some countries, such as the United States of America where E911 is deployed [50], it may not be possible for an LI subject to prevent its LI being used for this purpose.

- *Other services* that may be provided to entities other than the LI subject include LI subject tracking. This enables a range of possibilities, including allowing an employer to track employees to efficiently manage resources, and allowing a car-leasing company to track their cars.

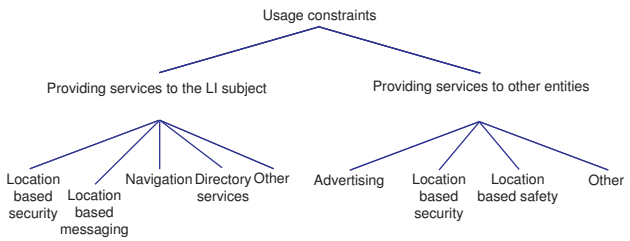

Figure 4.1: LI usage tree

If an LI subject wishes their LI to be used for one specific purpose, the use of constraints allows this to be made clear.

4.4 Limitations of Constraints

We next review some of the fundamental limitations of our notion of constraints.

4.4.1 Difficulties in Preventing and Detecting Constraint Abuse

Clearly, once an entity is in possession of LI, it is free to do with it as it wishes; in particular, it can ignore any associated constraints. LI is data, and the constraints which may be set on it do not physically prevent the receiving entity from misusing it. What adding constraints does do, however, is to allow entities to know the wishes of the LI subject. A regulatory authority which oversees the way in which

other entities handle constraints may go some way towards preventing constraint abuse.

Another problem which arises when considering the use of constraints is proving when they have been ignored. We have already established the difficulties of preventing the abuse of constraints with LI; it is also difficult to prove an entity has abused LI.

4.4.2 LI Constraint Predicaments

The aim of using constraints with LI is to enable an LI subject to dictate the distribution and use of the associated LI. Applying constraints to LI may, however, lead to further security considerations.

When a user applies constraints to LI, they give information which indicates how, or how not, to treat the LI. Although this information may be necessary to prevent the misuse of LI, applying constraints means that further information is divulged to the receiving entity. Two examples of this are now discussed.

4.4.2.1 Time Constraint Predicament

Two potential schemes for time constraints were mentioned in section 4.3.1.1. One made use of a validity period for the constraints. The other is where the LI is valid until a specified point in time.

The first scheme makes use of a time stamp which is added to the constraint. This allows a receiving entity to calculate the time at which a user was at the location shown by the LI. Of course, in most cases, the entities who are likely to receive LI are in all probability trusted by the LI subject, and so the fact that they know that a user was at a location at a particular time should not be a problem. The difference between this and the second scheme for specifying time constraints is that, in the latter case, a receiving entity is not informed precisely when the LI subject was at a particular location. The delay before the receiving entity obtains the LI may only

allow an approximate location of the LI subject to be calculated. Of course, in some cases, LI may be used in real time and, in such cases, the second scheme may be inadequate. An example of this latter situation is provided by a navigational service, where the location and movement of the LI subject must be calculated in order to provide the required information.

4.4.2.2 Distribution Constraint Predicament

Although distribution constraints are required to provide an entity with information about how it may distribute LI, providing distribution constraint information to entities can itself be a potential privacy violation. This information provides the receiver with information about entities for which the LI subject may or may not have a preference. Moreover, the receiving entity could use this information to pose as a preferred entity so that it can receive LI. In chapter 5, we discuss a mechanism where LI and constraints can be freely distributed in encrypted form, and released by a third party only if they are permitted by the constraints.

4.5 Combining Constraints with Auditability

Preventing misuse of LI is inevitably a complex task. For an entity to be able to use LI, they must have access to it. After an entity has seen the LI, they thereafter can use or misuse it as they please. Even when constraints are bound to the LI, an entity may choose to ignore them or decide not to pass them on.

Instead of trying to prevent misuse of LI, which is almost certainly an impossible task, we therefore propose the concept of auditability of LI. The idea is to enable all users of LI to determine where LI originates from, and to make all users accountable for their uses of LI. To work effectively, the majority of the users of LI must abide by the auditability rules, but this seems a reasonable assumption (otherwise there is little hope of achieving any control over LI). Of course, auditing will not prevent abuse, but it does enable misuse to be detected after the event, thereby acting as a deterrent to misuse.

The notion of auditability introduced here requires use of digital signatures. Every piece of LI, and its associated set of LI constraints, must be accompanied by a digital signature computed over both the LI and its constraints. That is, when any LI is generated by an LI gatherer, then, as well as generating and attaching the LI constraints, the LI gatherer must create a signature over the LI and the associated constraints. The LI gatherer might also be required to provide evidence with the LI of how the LI was obtained, and include this evidence within the scope of the signature.

Any entity receiving LI must verify the accompanying signature, and must log an exception (and must not use the LI) if the signature verification fails or if the signature is not present. Moreover, all LI users must check the constraints accompanying received LI to determine whether they should be in receipt of the LI — again, if they are not, then an exception should be generated and the LI should not be used. Finally, the LI and the signature should be retained for auditing purposes for a specified period of time. This may be dictated by the constraints themselves. In cases where this conflicts with legal requirements, the legal requirements must be considered first.

We now consider how this combination of rules can prevent (or at least make more difficult) the unauthorised distribution of LI. First observe that the mechanism described above does not address the misuse of LI, i.e. the use of LI in ways prohibited by the LI constraints. It is instead intended to address the issue of unauthorised distribution of LI (after all, uncontrolled dissemination of LI is probably the issue of greatest concern to most LI subjects). It can detect an entity that distributes LI in an unauthorised manner to an honest entity.

Suppose a malicious entity wishes to redistribute LI in a way prohibited by the LI constraints. If the entity simply sends it on as received, then the recipient will detect that the constraints have been violated and the malicious entity can be held responsible for the breach of constraints. Hence the malicious entity will need to change the LI constraints. This, however, will invalidate the original signature, and sending the LI without a signature will also enable the recipient to detect an LI use violation. Hence, if an entity wishes to disseminate LI with modified constraints, then they must sign the LI and indicate from where it was obtained — this may

present a major problem for a fraudulent LI user. It will, at minimum, enable a subsequent audit to detect exactly which entity was responsible for disseminating unauthorised LI.

A further measure to restrict the ability to fraudulently disseminate LI would be to limit the entities capable of acting as LI gatherers and generating signatures on LI. If an LI gatherer required a licence (e.g. in the form of an attribute certificate) to generate signed LI, then a malicious user without such a licence could not falsely disseminate LI, except to other malicious users.

Clearly this notion of auditability is dependent on industry co-operation and a regulatory body to ensure that rules are obeyed. However, whilst not universally effective, this is essentially the approach followed in many other spheres, including large parts of the financial industry.

4.6 Specification of Constraints

In order to enable a wide use of location-based services it is important to have a single language for the specification of LI. This should allow LI to be generated, transferred and used on a wide variety of platforms. Currently the most promising means of achieving a universally recognised means of specifying LI would be to employ an appropriately devised XML schema. XML (eXtensible Markup Language) is a language for data exchange between devices [19]. It allows data to be shared regardless of programming language or operating system, making it a strong candidate for use with location-based services and to describe LI. If LI is described in XML, it should also be possible to describe constraints in XML, giving similar advantages. XML can also be used to create digital signatures, which may be used to support the auditing scheme mentioned above.

4.7 Summary

Although attaching constraints has the advantage of allowing entities to see the requirements of the LI subject regarding LI, in so doing it also allows them to see additional information which may breach the privacy of the LI subject. It is also difficult to ensure that entities abide by the constraints which are set by the LI subject, and to prove when the constraints have been abused. Finding ways to address such issues is an important research challenge, and one that is addressed in the next chapter.

Location-based services are just one type of context-based service. A context-based service is one in which the context of an application automatically initiates some activity. LI in particular is an example of behavioural information, discussed further in Chapter 1. Certain intrinsic properties of such information mean that protection mechanisms must be designed specifically for it. For example, LI can be gathered from a number of different sources. Consequently, its distribution must also be controlled. This may require that constraints are available to a number of different entities.

Examples of context other than location include temperature and special events. Of course different forms of context have different security aspects. For example, the temperature of a subject's environment may not be private data; however, the end user's personal blood temperature may be private. As with LI, it would be necessary to subject such data to distribution and use constraints. This would mean extending the constraints described here to different contexts.

Although this chapter has identified some of the constraints which may be set by the LI subject, a formal language for specifying these constraints will enable their verifiable and unambiguous use. An investigation into current schemes such as P3P and PIDF for use in an LI-based environment, and defining such a language, may be a fruitful area for future research.

Chapter 5

The Location Information Preference Authority

Contents

5.1 Introduction **73**
5.2 A Mechanism to Provide Security for Constraints **74**
5.2.1 Overview of the Mechanism 74
5.2.2 Requirements for Use of the Mechanism 77
5.2.3 LI Token Creation 78
5.2.4 LI Distribution 79
5.2.5 LI Use 80
5.3 Billing **82**
5.4 Performance Analysis **83**
5.4.1 Assumptions 84
5.4.2 Storage Requirements 85
5.4.3 Message Exchanges 86
5.4.4 Operations Performed 87
5.5 Security Analysis **89**
5.6 Summary **91**

In Chapter 4 we discussed the fact that location-based service providers need to have access to LI regarding the users which they wish to serve, and that this is a potential privacy threat. We proposed the use of constraints, i.e. statements limiting the use and distribution of LI, that are securely bound to the LI, as a means to address this threat. We found that constraints may themselves reveal information to any potential LI user — that is, the constraints themselves may also be a privacy threat. To address this problem, this chapter introduces the notion of a LI Preference Authority (LIPA). A LIPA is a trusted party which can examine LI constraints and make decisions about LI distribution without revealing the constraints to the entity requesting the

LI. This is achieved by encrypting both the LI and the constraints, so that the LI is only revealed at the discretion of the LIPA. Much of the work described in this chapter has previously been published in [58, 59].

This chapter begins in section 5.1 with an introduction to the LIPA. Details of the LIPA model are described in section 5.2. The commercial aspects of the LIPA model are discussed in section 5.3. A performance analysis for use of the LIPA model is described in section 5.4. A security analysis of the model is given in section 5.5. Finally, the findings of the chapter are summarised in section 5.6.

5.1 Introduction

As previously discussed in this thesis, in order to offer location-based services, service providers need to have access to LI regarding the users they wish to serve; this is clearly a potential privacy threat. We have proposed the use of constraints, i.e. statements limiting the use and distribution of LI, that are securely bound to the LI, as a means to address this threat. As pointed out in section 4.4.2, constraints may themselves reveal information to any potential LI user — that is, the constraints themselves may also be a privacy threat. To address this problem, we introduce here the notion of a LI Preference Authority (LIPA). A LIPA is a trusted party which can examine LI constraints and make decisions about LI distribution without revealing the constraints to the entity requesting the LI. This is achieved by encrypting both the LI and the constraints with a LIPA encryption key, ensuring that the LI is only revealed at the discretion of the LIPA.

Other schemes that use a trusted party to convey data to enquirers include Single Sign On (SSO) systems. Such a system allows a user to authenticate themselves once to an entity, and then gain access to a number of services. A popular example of such a scheme is Kerberos [95]. This scheme provides authentication for a number of services in a distributed environment. In this scheme, a Key Distribution Centre (KDC) known as the Authentication Server (AS) acts as the trusted party. The user first authenticates itself to the AS, which provides it with a Ticket Granting Ticket (TGT). The TGT can then be used by the client to prove that it has previously been authenticated when requesting a service. A second KDC (called

the Ticket Granting Server) then generates a second ticket that allows the user to authenticate itself to the service.

The fact that constraints may themselves be regarded as personal information motivates the design of the scheme proposed in this chapter. The Location Information Preference Authority (LIPA) enables the end user to take advantage of location-based services, whilst controlling the way LI is used, stored and distributed. A LIPA is essentially a trusted party which helps control the distribution of LI and accompanying constraints.

We suppose that LI is distributed to service providers in the form of an LI token. The LI token includes LI securely bound to its constraints. The LI and constraints are also encrypted using the LIPA's private key, ensuring that unauthorised entities cannot see this information.

Finally, we discuss further work which may aid the wider use of multimedia location-based services.

5.2 A Mechanism to Provide Security for Constraints

In this section we introduce the LIPA mechanism that provides privacy control for LI and associated constraints.

5.2.1 Overview of the Mechanism

In order to ensure that the information held within the constraints remains private, we propose the use of a trusted party, which we call a Location Information Preference Authority (LIPA). The LIPA is responsible for deciding, based on given constraints, whether an LBS provider is allowed to have the LI of an LI subject. The information sent to the LIPA is encrypted in an LI token so that other entities cannot view it. This issue allows general distribution of LI within the scope of the LI token. The LI gatherer is assumed to be in possession of the list of preferred LIPAs for each LI subject for which it generates LI. This list will contain LIPAs

trusted by the LI subject. The LI gatherer must be trusted by the LI subject to act according to its wishes.

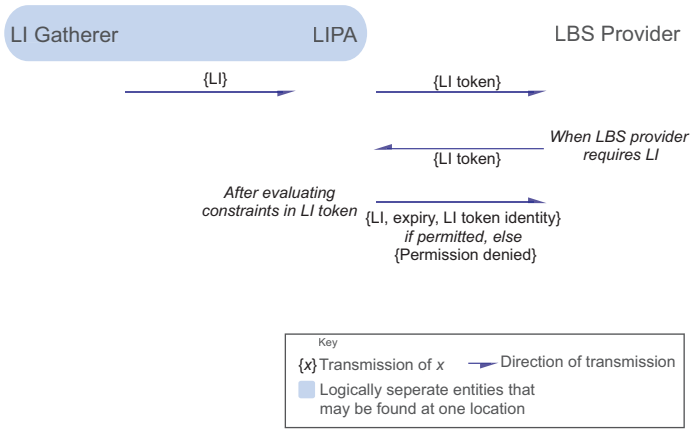

Figure 5.1: LIPA summary

Figure 5.1 gives a summary of how a LIPA can be used to relay LI using an LI token. This is discussed in greater detail below.

1. **LI gathering.** The first step in our mechanism involves the provision of LI by the gatherer. The LI gatherer may be at any location, including in the UD itself. The LI gatherer may obtain LI in response to a request by an LBS provider or an LI subject, or it may constantly collect LI for a large number of LI subjects.

2. **LI token generation.** The LI gatherer then creates what we refer to as an LI token. This includes both LI and accompanying constraints. The LI and constraints are encrypted by the LI gatherer using the public key of the LIPA. This ensures that only the LIPA is able to view this information. Also contained within the scope of the token is information which helps to identify both the LI subject and the LIPA, together with a unique token identifier. The LI token includes the signature of the LI gatherer, guaranteeing the integrity of the LI token. This also provides evidence to receiving entities regarding the identity of the LI gatherer. An LI gatherer may generate several tokens for the same LI, e.g. if an LI subject uses two or more LIPAs. There is also provision for the inclusion of an optional public key certificate for the LI gatherer's public key.

3. **LI token distribution.** When LI is required, an LI token is provided to the LBS provider wishing to use the LI for service provision. This could occur in a variety of ways, e.g. by using third party LI token repositories, by sending the LI token via the UD, or by direct transfer from the LI gatherer to the service provider.

4. **LI token verification and decryption.** Once an LBS provider wishing to use LI receives an LI token, it must submit it to the appropriate LIPA. From the LI token the LBS provider can establish the identity of the LI subject, the identifier for the LI token, and the identity of the LIPA, but not the LI or constraints, since they are encrypted.

   Upon receiving the LI token, the LIPA verifies the signature and then decrypts the LI and the constraints, and checks whether access to LI is permitted for the requesting LBS provider. If access to LI is permitted by the constraints, the LIPA returns the LI, the date/time of expiry of the LI, and the identifier of the LI token, all encrypted with the public key of the LBS provider and signed by the LIPA. If permission is denied, a message stating this is sent to the LBS provider. Note that, as discussed in section 5.4.3, the LIPA may choose only to send part of the LI, or a degraded version of it, to the LBS provider, if the constraints so specify. For example, the LBS provider may only be authorised to send LI accurate to at most 100 metres to this particular LBS provider, in which case the LIPA must reduce the precision of the LI before sending it.

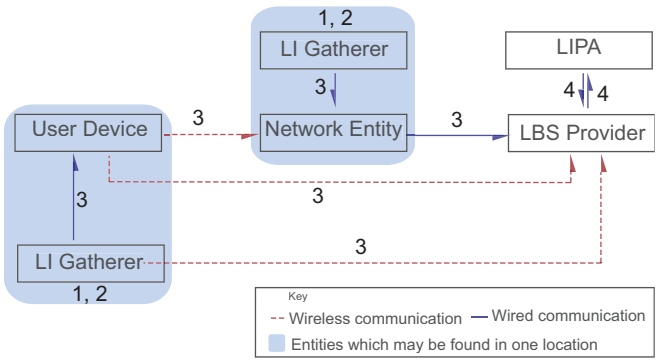

Figure 5.2: LI-related transmission for a mobile UD

Figure 5.2 shows the LI messages which may be transmitted in a mobile environment, where the numbers correspond to the four numbered paragraphs above. Stages 1

and 2 take place at the LI gatherer. Stage 3 indicates the distribution of the LI token. When entities are at the same location, e.g. when both are on a UD, wired transmission is assumed. In stage 4, the LI token is sent to the LIPA by the LBS provider, and LI may be provided in return.

5.2.2 Requirements for Use of the Mechanism

This section describes the requirements on the entities involved in use of the mechanism.

The LI gatherer is the entity responsible for creating LI. It must possess a signature key pair. It must also possess a trusted copy of the public encryption key for all the LIPAs used by the LI subjects for which it generates/collects LI. These keys are used to encrypt the LI and the constraints in the LI token. The LI gatherer must also be in possession of a reliable copy of the constraints and LIPA preferences for each LI subject for which it generates LI.

The LIPA must possess both a signature key pair and an asymmetric encryption key pair. It must also possess a trusted copy of the public verification key of every LI gatherer whose LI it needs to process, and a trusted copy of the public encryption key of each service provider to whom it might wish to provide decrypted LI. (The need for LIPAs to hold public keys of LI gatherers and LBS providers can be obviated by requiring LI gatherers and LBS providers to obtain and distribute public key certificates).

Each LBS provider must possess a trusted copy of the public signature verification key of each LIPA with which it interacts. It must also possess an asymmetric encryption key pair. As one can imagine, this may pose a significant key management task. We envisage that the LIPA may also handle key management tasks. Entities that wish to verify a LIPA's signature can obtain a public verification key from that LIPA when required. The LIPA will also need to possess the public encryption keys of all the LBSs that register with it. Such a public key may be provided when an LBS registers itself with the LIPA. This registration may also require that the LBS divulges how it intends to use LI that it will request in the future.

It is assumed that all the necessary encryption and signature algorithms have been globally agreed before use of the scheme.

5.2.3 LI Token Creation

The entity responsible for generating LI is also responsible for creating what we refer to as an LI token. At the time of creation (or acquisition) of the LI, we suppose that the LI gatherer generates accompanying constraints C based on pre-specified LI subject preferences. The structure of the LI token is described below.

LI Token:

$E_{e_L}(LI\|C)\|$
$I_L\|I_U\|TokenID\|I_G\|$
$S_G(E_{e_L}(LI\|C)\|I_L\|I_U\|TokenID\|I_G)\|$
$[Cert_G]$

where:

e_X represents the public encryption key of entity X; E_{e_X} denotes asymmetric encryption using the public key e_X; $X\|Y$ represents the concatenation of data items X and Y; L represents the LIPA; U represents the LI subject; G represents the LI gatherer; I_X represents an identifier for entity X, e.g. I_G denotes an identifier for the LI gatherer G; $Cert_G$ is the public key certificate of the LI gatherer; and [...] represents an optional data item.

The LI token is divided into four parts: the encrypted part, the plaintext part, the digital signature, and the (optional) public key certificate of the LI gatherer.

- The encrypted section consists of an encrypted version of LI and the constraints, C. These are encrypted using the public key of the LIPA, e_L. This ensures that entities other than the LIPA cannot see this information.
- The plaintext part consists of I_L, I_S, $TokenID$ and I_G. The identifier I_L identifies the LIPA whose public key has been used to encrypt the LI and the

constraints. This enables any entity wishing to gain access to the contents of an LI token to determine which LIPA it can be requested from. This identifier could take a variety of forms, e.g. a URL or an IP address. The identifier I_U allows any entity to identify the LI Subject to which the LI in the token relates. This identifier may be a pseudonym. The $TokenID$ is an identifier which, in conjunction with I_G, enables an LI token to be uniquely identified. The identifier I_G allows any entity to determine which entity generated the LI token. This also enables entities to decide which public key to use to verify the digital signature. This identifier may also be a pseudonym.

- The digital signature is computed over both the encrypted and plaintext parts of the LI token. This provides assurance that the LI Token has not been tampered with, and authenticates the entity which created the LI.
- The certificate $Cert_G$ may optionally be included in the LI token. This makes it easier for LIPAs which communicate with many LI subjects to obtain the necessary public keys.

Before proceeding, note that the encrypted part of the LI token could alternatively be encrypted using a symmetric encryption scheme with a shared secret key. The major advantage of such an approach would be that a symmetric encryption algorithm is typically much less computationally intensive than an asymmetric scheme. The main disadvantage is the key management overhead, since such an approach would require each LI gatherer to share a secret key with every LIPA with which it 'does business'. A variety of different mechanisms exist to provide the necessary key management functions — see, for example, [95].

5.2.4 LI Distribution

Section 5.2.3 describes the structure of an LI token. When there is a request for LI, or when an LI subject requests a service, the LI token is sent to the relevant LBS provider.

LI Gatherer $\rightarrow P$:

$E_{e_L}(LI \| C) \|$

$I_L \| I_U \| TokenID \| I_G \|$
$S_G(E_{e_L}(LI \| C) \| I_L \| I_U \| TokenID \| I_G) \|$
$[Cert_G]$

where:

$A \rightarrow B$ represents the communication of a message from entity A to entity B;
P represents the LBS provider.

LI should always be distributed within an LI token, regardless of who is sending the LI. The message above describes direct communication of the LI token from the LI gatherer to the LBS provider; however, as mentioned earlier, LI tokens may also be distributed via third parties and between LBS providers.

5.2.5 LI Use

This section describes how an entity uses an LI token. When a LBS provider wishes to gain access to the LI within an LI token, it must send the LI token to the LIPA whose identifier is in the token, and hence whose public key was used to encrypt the LI in the token.

$P \rightarrow$ LIPA entity:
$E_{e_L}(LI \| C) \|$
$I_L \| I_U \| TokenID \| I_G \|$
$S_G(E_{e_L}(LI \| C) \| I_L \| I_U \| TokenID \| I_G) \|$
$[Cert_G] \| [Cert_P]$

The above indicates the LBS provider sending the LI token to the LIPA entity. The LBS provider may optionally include a certificate for its public key, to avoid the need for the LIPA to possess a trusted copy of every LBS provider's public key. When the LIPA receives the LI token, it must first verify the signature and then decrypt the enclosed LI and constraints. If the signature is invalid, or the token syntax is not as expected, then the LBS provider must be sent the 'Permission Denied' message

(see below). The LIPA must then check that the LBS provider is permitted by the constraints of the LI subject to receive this LI. The LIPA must also check the authenticity of the LBS provider, which may be based on the certificate provided by the LBS provider. Details of a mechanism to provide this check for authenticity are not discussed further here.

If the LBS provider is permitted to have LI, then the LI will be sent by the LIPA to the LBS provider. The structure of the message used to send the LI back to P is described below. The LIPA also keeps a record of the LI token and the entity to which it is providing LI.

LIPA entity $\rightarrow P$:

$E_{e_P}(LI \| Expiry \| TokenID)$

$S_L(E_{e_P}(LI \| Expiry \| TokenID))$

The message from the LIPA to the service provider contains two parts: the encrypted part which contains LI, $Expiry$ and the $TokenID$, and the signature. The encrypted part is encrypted with the public key of the service provider requesting the LI. This ensures that only this service provider can read the information, preventing malicious parties intercepting data while in transit. $Expiry$ is a time-stamp extracted from the constraints, and specifies when the LI expires, i.e. when the LI should be deleted by the service provider. This is the only information from the constraints which needs to be sent to the service provider. The $TokenID$ allows the LI subject to relate the LI received from the LIPA to the LI token from which it has been taken. The digital signature allows the receiving entity to check whether the message has been tampered with during transit.

If the requesting entity is not permitted to have access to the LI in the token, then the following $PermissionDenied$ message is sent to the requesting entity:

LIPA entity $\rightarrow P$:

$TokenID \| PermissionDenied$

5.3 Billing

There are two main approaches to billing in the LIPA model. First, the entity that receives the service could pay the billing entity every time that it utilises a service. As demonstrated in the example scenario in section 4.2.1, a user may use many services in a short space of time. This may be inconvenient, because the user has to pay for each service as they use it. This leads us to the second approach, in which the entity that utilises the service is billed for the services periodically. This approach typically requires that the billing entity keeps a record of the services that the user has utilised. In such cases, the privacy requirements, in particular the storage time constraints for the LI subject, may conflict with the billing requirements. However, the LI subject will typically have a contractual agreement with the billing entity, and would be aware of the requirements of that entity before agreeing to the service. This should not affect the usage and distribution constraints.

There are numerous ways in which the LIPA could generate income for the provision of its service. The LIPA may charge for each request for LI which it receives, or each successful request for LI, i.e. when LI is sent to a LBS provider by a LIPA. Also, billing may be per LI token or per individual request.

The entities which could be billed for the LIPA service are the LI subject and the LBS provider. Billing the LI subject may result in a scenario where LBS providers could request LI from the LIPA, which will charge the LI subject whether or not the LBS provider gives any service to the subject, and this is clearly not desirable. Alternatively, billing the LBS provider appears a more appropriate solution, since the LBS provider can recover the cost of obtaining the LI from the entities to which it provides service.

The LI gatherer (unless it is the LI subject him/herself) will also typically require a means of obtaining payment for providing LI tokens. The LI gatherer could charge the LI subject or the LBS provider when it is providing an LI token to it directly. When the LBS provider is charged, they can recover the cost from the LI subject in providing the service.

5.4 Performance Analysis

For the purposes of the analysis in this section, we first imagine a scenario where an LI subject wishes to utilise an LBS. We assume that the LI subject possesses a GPS-enabled mobile telephone that accesses a 3G network [94]. This device acts as the UD in our model. The GPS device has the role of the LI gatherer. This will be responsible for generating the LI token. The LI in the token will be GPS coordinates, that are converted to the 3G format. Location Information in this context is known as LOCI [3]. Using the LI token provided by the LI gatherer, an LBS provider could then send the LI token to the LIPA. The size of the data that is described in this section is based on this scenario.

Other types of LI and an estimate of their respective lengths are shown in Table 5.1. Degrees, minutes and seconds are the most commonly used representation of geographic coordinates. In this system, two numerical values are used to represent latitude and longitude. This data can also be represented in decimal degrees, where the minutes and the seconds are converted into decimal values. GPS data is represented in decimal degrees.

IP addresses are assigned to network devices in IP networks allowing them to communicate with other networked devices. Such addresses are usually assigned on a country-wide basis. Additionally, IP address assignation is publicly available, making it possible to locate network devices. The address size used by IPv4 is 32 bits. This has increased to 128 bits in the specifications for IPv6.

The Universal Transverse Mercator (UTM) system specifies location using a grid-based system. It divides the earth into 60 different zones. Latitude and longitude coordinates are then used within these zones to specify a location. UTM data is represented in 3-tuples; the zone, the latitude, and the longitude.

The scenario for our analysis involves an LI subject with a wireless UD who wants to receive a service based on his or her location. We assume the existence of a wireless network in which the LI gatherer is a network entity. The UD and the LI gatherer are capable of both transmitting and receiving wireless transmissions.

Table 5.1: LI Lengths

	LI Length (bits)
Decimal degrees	64
IPv4 address	32
IPv6 address	128
UTM	100
3G LI (LOCI)	120

5.4.1 Assumptions

We suppose that encryption of LI token contents is achieved using a 'digital enveloping' technique. This involves first generating a random secret key, K say, which is used to encrypt the data (using a symmetric encryption algorithm). The secret key K is then itself encrypted with an asymmetric encryption algorithm (using the public key of the intended recipient), and sent with the encrypted data. For the purposes of our analysis here we suppose that K contains 128 bits, that symmetric encryption takes place using a stream cipher, e.g. AES in counter mode, and that asymmetric encryption uses 1024-bit RSA (e.g. using OAEP) (see Chapter 2). We moreover suppose that the public encryption exponent is always 16 bits long, i.e. the total length of a public encryption key is 1040 bits (see Chapter 2). Hence, since we are supposing that symmetric encryption leaves the length of the data unaltered, enciphered data strings will always contain 1024 bits more than the corresponding plaintext strings.

The digital signature scheme used is also based on RSA with a 1024-bit key. To generate the RSA digital signature, the data is first hashed using the SHA-1 hash algorithm. This gives a 160-bit output (see Chapter 2 for further information).

The LI gatherer is capable of generating RSA digital signatures, and performing RSA encryption. It is also capable of generating a hash using the SHA-1 algorithm. The LBS provider is capable of verifying digital signatures and decrypting data. The LIPA is capable of verifying and generating digital signatures, and also encrypting and decrypting data. The LI gatherer and the LIPA can also generate public and private key pairs for digital signatures. The LBS provider and the LIPA should be able to generate a public and private key pair for encryption.

The length of the LI is assumed to be 15 bytes, where 1 byte is 8 bits. This is a reasonable assumption, as the LI in a 3G network, called the LOCI, consists of 11 bytes. The LI subject identifier, I_U, is assumed to contain 10 bytes. The International Mobile Subscriber Identifier (IMSI) [3] in a 3G network, used to uniquely identify mobile subscribers, contains 9 bytes, so this is also a reasonable assumption. The LIPA identity I_L, the token identity $TokenID$, and the LI gatherer identity, I_G are all assumed to be 5 bytes in length, which allows a large number of unique identifiers. The constraints are assumed to be 1600 bytes long. This is the approximate size of an XML document consisting of approximately 40 lines. As we see below, this means that the total size of the LI token in this scenario is 1896 bytes.

5.4.2 Storage Requirements

In this scenario, we assume that LI tokens are generated upon request. This means that the LI gatherer is not required to store LI tokens. The LI gatherer must, however, obtain the 1040-bit public keys of LIPA entities to which it sends LI tokens. These may be stored by the LI gatherer or acquired when necessary. The LI gatherer must hold the constraints for each LI subject to which it provides a service. This is necessary to generate the LI token.

When the LBS provider receives the LI token it can decide whether it wishes to obtain the LI for the LI subject, based on its stored policies. These policies may list the LI subjects which have subscribed to a service; they may also specify the LI subjects to which a service should not be provided. If the LBS provider successfully receives LI from the LIPA, it may also store the LI until the specified expiry time.

The LIPA may also store policy information. This could include information about the LI subjects to which it provides service. These stored policies may also hold information about the activities of various LBS providers. This helps the LIPA's decision-making process when deciding if it should send LI to an LBS provider.

Storage requirements are summarised in Table 5.3.

5.4.3 Message Exchanges

The processes and data transfer requirements are summarised below and in Table 5.2. We start by describing the LI gatherer. After receiving a request for LI, it must perform the following tasks.

1. Encrypt the 1615 bytes of data (1600 bytes for the constraints and the 15 bytes of location data) using the RSA-based digital enveloping technique with the public key of the LIPA. This will result in an encrypted block of data containing 1615+128=1743 bytes.
2. I_L, I_U, $TokenID$ and I_G, i.e. a total of 25 bytes of data (10 bytes for I_U and 5 bytes for each of I_L, $TokenID$ and I_G), are then concatenated with the encrypted data described above. This will result in a data string containing 1768 bytes.
3. Generate the SHA-1 hash of the encrypted data and the identifiers, and then generate a digital signature over this hash using the private signature key of the LI gatherer. The resulting signature will contain 128 bytes. The concatenation of this with the identifiers and the encrypted data will result in a 1896-byte LI token.

The LI token is then sent to the LBS provider. When the LBS provider receives the LI token it decides whether it requires LI from this LI token. If this is required it then sends the 1896-byte LI token to the LIPA.

At the LIPA:

1. The LIPA first verifies the signature contained in the LI token.
2. The ciphertext in the LI token is then decrypted. Based on the constraints found in the resulting plaintext, the LIPA then decides whether the requesting LBS provider is permitted to receive the LI. If transfer of the LI to the service provider is permitted, it may be necessary to degrade the accuracy of the LI, depending on the constraints (see section 4.3).

Table 5.2: Message Exchanges

From \ To	LI Gatherer	LBS Provider	LIPA
LI Gatherer	N/A	1896-byte LI Token	1896-byte LI Token
LBS Provider	Y	N/A	1896-byte LI Token
LIPA	N/A	876-byte version of Token	N/A

Table 5.3: Storage Requirements

	Key Storage	Constraint Storage	LI Token Storage
LI Gatherer	LIPA public key	LI Token	LI Token
LBS Provider	Y	N/A	LI Token
LIPA	N/A	Small Token	N/A

3. The LIPA must then encrypt the data to be sent to LBS provider. The data to be encrypted in this case will be the 15 bytes of LI, and approximately 600 bytes for the constraints, as the expiry time will be the only data from the constraints to be sent[5.1]. This is sent together with the 5-byte $TokenID$. These 620 bytes of data will need to be encrypted, giving 748 bytes of ciphertext.

4. The above 748 bytes are then signed, resulting in a 128-byte signature. The resulting 876 bytes are then sent to the LBS provider.

5.4.4 Operations Performed

Tables 5.4– 5.7 describe the number of signature generations, signature verifications, encryptions, decryptions and certificate verifications that are required at various stages in the LIPA model. We describe two different scenarios at each stage. Scenarios 1 and 2 respectively cover the cases where symmetric and asymmetric encryption is used. For both scenarios we describe the computations required for LI token creation, LI token verification (by the LIPA), transmission of LI by the LIPA to the LBS provider, and receipt of LI by the LBS provider.

Table 5.4 shows the operations that are performed when the LI token is generated. To create the token, the LI gatherer first encrypts the LI and the constraints. As described in table 5.1, the LI can range from 64 to 128 bytes of data. The constraints

[5.1] This estimate for the data size of the constraints with only expiry time information is based on an XML document consisting of approximately 16 lines.

Table 5.4: LI Token Generation Operations

Scenario	Encryption	Decryption	Signature Creation	Signature Verification	Certification Verification
1	1	0	1	0	0
2	1	0	1	0	1

Table 5.5: LIPA Operations When LI Token is Received

Scenario	Encryption	Decryption	Signature Creation	Signature Verification	Certification Verification
1	0	1	0	1	1
2	0	1	0	1	1

can be in the region of 1600 bytes in length. In scenario 1 the secret key shared by the LI Gatherer and the LIPA must be located and used for the encryption. In scenario 2, a public key certificate for the LIPA must be obtained and verified. The public key must then be used to encrypt the LI and the constraints. There is also one signature creation during this process.

Table 5.5 shows the operations that are performed when the LIPA receives an LI token. The LIPA entity must first verify the signature in the LI token. If it has not been included in the scope of the LI token, a public key certificate must be retrieved. This must also be validated. If the signature is valid, the LIPA entity can then decrypt the LI and constraints. Scenario 1 requires that the shared secret decryption key is located and used to decrypt the data. Scenario 2 requires that the LIPAs private key is located and used to perform the decryption. If the LBS provider is permitted to have this LI, the LIPA must perform the operations described in Table 5.6. This process requires that the encryption key for the LBS provider is first located. This encryption key is then used to encrypt the LI, expiry time, and the token identifier. Scenario 1 requires that the shared secret key is located. Scenario 2 requires that a certificate for the LBS provider is retrieved and verified. In both cases the data is then signed and sent to the LBS provider.

Table 5.6: LIPA Operations When LI is Sent

Scenario	Encryption	Decryption	Signature Creation	Signature Verification	Certification Verification
1	1	0	0	1	0
2	1	0	0	1	1

Table 5.7: LBS Provider Operations

Scenario	Encryption	Decryption	Signature Creation	Signature Verification	Certification Verification
1	0	1	0	1	0
2	0	1	0	1	1

Table 5.7 shows the operations performed when LI is received by the LBS provider. The signature must be verified in both scenarios 1 and 2. A certificate for the public key necessary to verify this signature must first be retrieved and validated. This can then be used to verify the signature. The LBS provider must then retrieve its own decryption key to retrieve the LI, expiry time, and the token identifier.

5.5 Security Analysis

In this section we describe how our mechanism addresses control and privacy issues for LI. We refer to the threats to LI previously discussed in Chapter 3.

The primary aim of this chapter was to provide a mechanism that allows an LI subject to control access to LI using constraints. A secondary aim was to control access to the constraints themselves.

Enabling the LIPA to make decisions based on constraints ensures that entities that are not trustworthy cannot gain access to the LI or constraints.

This chapter has shown that entities can be in possession of an LI token, where the information that is held within it is only available to the LIPA and authorised LBS providers. The threats discussed in Chapter 3 provides a basis for the analysis in the rest of this chapter.

- **Unauthorised LI Use.** Unauthorised entities that are in possession of LI may also it in an unauthorised way. Of course, authorised entities may also use it in an unauthorised way; however, there is assumed to be a trust relationship between these entities and the LI subject. We have shown how the distribution of LI can be limited to those that are authorised to possess it in section 5.2.5.

We limit this by applying distribution, time and use constraints.

If an entity wishes to redistribute the LI of an LI subject, it should only distribute the LI token. If it chooses to redistribute LI in other forms, then this can only be addressed by some form of policing, e.g. through peer enforcement. Of course this protection could be enhanced by a regulatory authority, which would ensure that rules are being adhered. Chapter 6 investigates a means to provide further assurances about such entities.

A threat to LI use arises if an authorised entity is compromised by a malicious party. The authorised party could be an LI gatherer, a LIPA, or an authorised LBS provider. When an LI gatherer becomes compromised, it can potentially provide LI to malicious entities that can use it to track an LI subject. For example, malicious software on a UD could initiate such a compromise. In such cases, users should limit the likelihood of such an attack by only installing trusted software. When the LI gatherer is a separate entity, the risk of compromise must be reduced. Compromised LIPAs and LBS providers can similarly allow a malicious party to access LI, also potentially leading to its unauthorised use.

- **Eavesdropping Threat.** When LI is required, it must be sent within the scope of an LI token. This ensures that when it is in transit, eavesdroppers are not able to access the LI or constraints. If an entity is permitted to access LI, the LI, token identifier and expiry time must be encrypted and sent. This ensures that an eavesdropper cannot access any information that is being sent.

  Of course, the identifiers for the LIPA, the LI subject, and the LI gatherer are sent unencrypted, and this information will be available to an eavesdropper. However, this information cannot be used to learn about the LI and constraints that are within the token.

- **Unauthorised LI Modification.** The proposed scheme also addresses threats from unauthorised modification, through the inclusion of a signature found in the LI token. When an entity receives an LI token, it can verify the signature, proving that the LI token has not been modified, and also providing evidence of the entity that first created the LI token.

As discussed previously, once an entity is in possession of LI, maintaining control of

this information is a difficult task. Ensuring that LI is managed according to the preferences of the LI subject once an entity possesses it, can only be based on trust. Our mechanism aims to provide LI only to entities which can be trusted, giving the LI subject control over their LI. Of course, even trusted entities cannot be trusted all the time. Once these trusted entities have the LI, the LI subject can only rely on a regulatory or legal authority to ensure that messages are being transmitted in the manner which has been previously agreed.

Auditability should allow the identification of entities acting in violation of the rules set by the constraints. To prevent unauthorised distribution of LI, its origin, i.e. the entity responsible for generating the LI token, must be verifiable. In addition, users of LI must be accountable for its use. Therefore, as discussed in chapter 4, if a malicious entity redistributes LI in a way prohibited by the LI constraints, the recipient will detect this, and the malicious entity can be held responsible for the breach of constraints.

5.6 Summary

This chapter addresses the issue of control and privacy of LI and associated usage constraints by introducing a trusted third party based framework. We have introduced a mechanism which gives the end user the ability to control their LI without having to divulge additional personal data.

Chapter 6

Using Trusted Computing to Enhance Location Privacy

Contents

6.1 Introduction **93**
6.2 Trusted Computing **94**
6.3 Next-Generation Secure Computing Base **96**
6.4 Personal Information Model **97**
6.5 Scenarios **98**
6.5.1 Registration Scenario 98
6.5.2 Location-Based Service Scenario 99
6.5.3 Medical Records 99
6.6 TCG Mechanisms **100**
6.6.1 TPM Identities 100
6.6.2 TCG Measuring, Reporting and Storing Processes 102
6.6.3 Sealing Data 103
6.7 Protecting Personal Information Using Trusted Computing **104**
6.7.1 Overview 104
6.7.2 Practicality 106
6.7.3 Using Trusting Computing with PI 107
6.7.4 Constraints, LIPA and LI Tokens 109
6.8 Other Approaches to Privacy Protection Using Trusted Computing **110**
6.9 Analysis **110**
6.10 Conclusions **111**

In this chapter we show how trusted computing can be used to enhance privacy for personal information, and in particular for location information. We will focus on how the mechanisms in the Trusted Computing Group (TCG) specifications can be used to enable the holder of personal information to verify the trustworthiness of a

remote host before transferring the personal information to that remote device. By so doing, greater assurance can be provided to end users that their expressed preferences for the handling of personal information will be respected.

Some of the work described in this chapter has been published in [56]. This chapter begins in section 6.1 with an introduction. Trusted computing platforms and the TCG are discussed in section 6.2. Section 6.3 describes Microsoft's Next Generation Secure Computing Base proposals, which build upon functionality specified by the TCG standards. A model for the handling of personal information is described in section 6.4. Various scenarios for the use of personal information are described in section 6.5. The mechanisms from the TCG standard that are of relevance to this chapter are described in section 6.6. Section 6.7 discusses how personal information can be protected using trusted computing. Other aspects of privacy that may be tackled using trusted computing mechanisms are discussed in section 6.8. Finally, the ideas proposed in this chapter are analysed in section 6.9, and the chapter concludes in section 6.10.

6.1 Introduction

In Chapter 5, a trusted party, referred to as a Location Information Preference Authority (LIPA), was introduced as a means to help protect the privacy of LI. A simple model for LI protection was introduced, in which LI is distributed within the scope of an LI token. This LI token has the property that only the LIPA entity is able to view the information held within it. Along with the LI, constraints are also included in the LI token. As discussed in Chapter 4, constraints are simply statements which limit the use, distribution and storage of LI.

In our simple model, a Location-Based Service (LBS) provider wishing to use LI may only view it by making requests to the LIPA. The LIPA makes a decision whether or not to send LI based on the constraints in the LI token. If permitted, the LI will then be sent to the requesting LBS provider. The problem remains that a dishonest LBS provider may either use the LI in unauthorised ways or may redistribute it to other service providers. Whilst this problem is clearly a very difficult one to deal with, in this chapter we consider ways in which trusted computing facilities might be

used to help address it. Rather than look at LI specifically, we look more generally at Personal Information (PI).

First note that service providers will almost inevitably be using a specific application running on one or more servers to manage and use PI. Verifying the integrity and provenance of this application software, as well as of the platform on which the application is running, will potentially give the user additional assurance that their data is managed in the proper manner.

In this chapter we explore how, using the mechanisms found within the TCG specifications [140, 141, 142], a device wishing to send private data to another device can reliably discover the software state at the destination device. This can then be used to help decide whether or not the data should be sent. Further to this we discuss how an entity is able to ensure that only software which is implicitly trusted by the sending entity is able to use this private data. Finally we show how our LIPA model may be extended to use these mechanisms to enhance the privacy of LI.

6.2 Trusted Computing

Physical access to a machine will typically allow the software integrity of that machine to be compromised. If a (secure) computer digitally signs a message, then trust in the message depends on trust in both the computer software that computed the signature, and the physical security of the underlying hardware (and in the correct application of security procedures by administrators). This makes sense in a conventional 'computer centre'. However, PCs are typically not stored in a physically secure environment; even though modern versions of Microsoft Windows (and Linux) have multi-user security features, users and programs often run as the 'administrator' user, and there are many ways that the operating system integrity can be damaged. Thus today, neither the user of a PC nor a communicating party can trust very much about a PC, despite major efforts to improve the security of operating systems. Anyone with access to the PC hardware can modify the operating system (e.g. by removing the hard disk and changing files).

Even if the user looks after the physical security of their PC, there are many other

threats to system integrity. Modern operating systems and applications are highly complex, and it is almost impossible to remove all vulnerabilities; moreover, users can accidentally run malicious software which can damage system integrity.

However, the user may nevertheless want to trust the integrity of their PC. For example, the PC may be used for managing a bank account, performing e-commerce transactions, or managing personal information, all of which require user trust in the PC. A third party may also want to trust the integrity of a PC. This could be for a variety of reasons, e.g. the third party is a bank and the PC is being used for e-commerce, the third party is a content provider and the PC is performing Digital Rights Management, or the PC is performing other security functions (e.g. authentication, key management) on behalf of a third party, all of which require third party trust in the PC. Trusted computing enables trust in the integrity of a PC based on a combination of software and hardware features. In particular trusted computing enables remote third parties to measure the integrity of the PC software environment. Detailed discussions of trusted computing can be found in [10] and [97]; see also [51] and [147].

In the past, a number of attempts have been made to realise trusted computing concepts in one form or another. Tygar and Yee [145] discuss the use of a secure coprocessor which ensures the integrity of the bootstrapping process and also system software. The secure coprocessor may also be used to verify data integrity in the same manner in which it ensures the integrity of the bootstrapping process.

Clark and Hoffman's Boot Integrity Token System (BITS) [29] involves the use of a smart card to provide an access control mechanism to preserve boot integrity. In addition, the smart card contains integrity values for executables found on host systems, which can be used for virus detection.

The AEGIS [7] architecture involves an integrity chaining process which is initiated at the bootstrapping process. This ensures that the boot process ends in a secure state. Only trusted or verified code may execute. If the code integrity is compromised, the architecture is designed so that it can still boot in a verified manner using an alternative module.

The Trusted Computing Platform Alliance (TCPA) was formed in 1999 to develop hardware specifications for trusted computing platforms. The original contributing members of the organisation included HP, IBM, Intel and Microsoft. The TCPA was superseded by the TCG in 2003. The aims of the TCG include the further development of the specifications produced by the TCPA. The notion of a trusted platform as defined by the TCG is discussed further in section 6.6. The TCG is also developing profiles of the specifications directed at a variety of types of host, including Personal Digital Assistants (PDAs), mobile telephones, and servers. The features of the TCG specifications used in this chapter are discussed further in section 6.6.

Numerous applications for systems conforming to the TCG specifications have been suggested, and we briefly mention three of these. Balacheff et al. [11] discuss the use of a trusted display controller to provide a trusted method to digitally sign documents using a smart card in an open computing platform. Chen et al. [25] discuss the application of trusted computing platforms for a biometric authentication system. Finally, Pashalidis and Mitchell [103] consider the use of trusted platforms to achieve Single Sign On capabilities. These and other applications are discussed in more detail in [97].

6.3 Next-Generation Secure Computing Base

Microsoft's Next-Generation Secure Computing Base (NGSCB) has been designed to take account of the fact that standard operating systems are large and complex [46]. To provide security and system integrity for such a system, it partitions the platform into two or more isolated operating systems. These include a secure operating system and a standard operating system. The two (or more) operating systems work concurrently, and their separation is guaranteed by an underlying isolation kernel. The isolation kernel ensures that the operating systems are isolated from each other. This feature can enable security essential processes to run in a different area from other processes running in the standard operating environment. Users may then choose between use of the secure environment in which they can take advantage of the secure mechanisms offered, and the standard operating system environment.

NGSCB provides various security services such as access control and integrity; however, services such as physical data storage take place in the standard operating system. Initialisation of the secure operating system relies on initalisation of the isolation kernel. The isolation kernel then partitions memory space. Security mechanisms found in the secure operating system environment include: process isolation, sealed storage, attestation and trusted paths.

- Process isolation ensures that guest operating systems and the area of the platform responsible for executing security essential processs are completely isolated.
- Sealed storage provides confidentiality and integrity for stored data. It also ensures that data is only accessible when the platform is in a state defined by the user and the secure operating system. This feature can also be used to ensure that data cannot be viewed when migrated to different platforms or operating systems. Code identities [46] indicate the platform state that must hold before data can be unsealed.
- Trusted paths to a user ensure that devices such as a keyboard, mouse and graphics interact with programs in a predictable manner.

Most importantly, the ideas presented in this chapter (and discussed in the context of the TCG specifications) may also be used in an NGSCB environment. Further information on NGSCB can be found in [46, 47, 104, 105].

6.4 Personal Information Model

We next describe a general model for the generation and use of personal information. We will use trusted computing to protect personal information in the context of this model.

- **Personal Information (PI).** This is data which provides personal information regarding the subject. PI may occur in many forms, ranging from telephone numbers to medical records.

- **Service Provider (SP).** This entity makes use of PI in the provision of services.
- **Subject.** The subject is the entity about whom personal information is gathered and used.
- **Trusted Platform (TP).** A trusted platform is one that contains a trusted subsystem which is able to both measure the current software state and reliably attest to this state to remote third parties (see Chapter 3). In the context of this chapter, this is the platform owned by the SP.
- **Regulator/Legal authority.** This is an entity which exerts legal or regulatory control over the management and use of PI. This includes telecommunications regulators, data privacy authorities, law enforcement bodies, and auditors.

Note that this model is based on a simplified version of the LI model introduced in section 4.2.

6.5 Scenarios

Using our model, we look at various scenarios involving the use of PI, and identify some of the associated risks.

6.5.1 Registration Scenario

Suppose a subject wishes to make a purchase using a website. Payment and delivery information such as credit card and address details will typically need to be provided for such a transaction to take place. By using a secure means of transmitting data, the subject can be assured that the confidentiality of data is preserved during transmission. One major concern which remains is the nature of the security mechanisms implemented by the service provider to protect its stored data. That is, even though the subject is assured that data was transmitted over a secure channel, in

many cases there is little or no assurance that data will be stored securely once it reaches the SP's server.

One example of an issue arising in this scenario is provided by those web sites that require a verified e-mail address in order to gain access to the services that they provide. In some cases these e-mail addresses have been sold to advertisers who send unsolicited mail to the subject. It would be desirable if a subject could ensure that e-mail addresses (and any other PI) are handled in a way consistent with the subject's privacy requirements.

6.5.2 Location-Based Service Scenario

Consider a subject that wishes to take advantage of a location-based service. The subject must provide its location to the SP (or allow the SP to obtain its LI from a third party) in order to enable the provision of such a service. This LI is clearly just a special type of PI. Once the LI has been provided to the SP the subject has no direct way to control how this PI will be used.

6.5.3 Medical Records

In our third scenario we consider a clinic where all personal medical records are held on a computer. Precautions should be taken to ensure that, when they are sent to other locations, records are not accessed by people who do not have the necessary access privileges. For example, if a patient is being treated at a specialist clinic at a different location, it is important that only the relevant practitioners have access to that patient's medical records. If some evidence of the conditions in which these records will be kept was available to the holder of the records, then it may aid the holder of the records in deciding whether or not to pass the records to this remote entity. It is also common for medical records to be accessed for research purposes. In such a case it would again be desirable if the clinic could establish the conditions in which medical records would be held by the researchers.

6.6 TCG Mechanisms

We now give a brief overview of the TCG mechanisms relevant to this chapter. Many of the mechanisms discussed here are described in detail in [60].

The TCG specifications [140, 141, 142], define a Trusted Platform (TP) as a computing platform which contains a Trusted Platform Subsystem (TPS). The subsystem contains the parts of the TP which provide fundamental trust and security capabilities. The mechanisms which seal, measure, store and report integrity metrics in a trusted manner are now discussed. Using a combination of these mechanisms we will subsequently show how it may be possible to ensure that data is being sent to a trusted machine.

The TPS contains the Trusted Platform Module (TPM), which is the Core Root of Trust for Measurement (CRTM), and the Trusted Software Stack (TSS). The TPM is an integrated circuit responsible for various security aspects of the TP. The TSS may be implemented as software and may not be trusted itself; however, it may be required for the operation of the TP.

When a subject communicates with a TP, it can decide whether or not to divulge information to it, depending on the state of the TP. Furthermore, the user may also ensure that the data it provides will only be accessible to the remote platform when that platform is in the same state as it is at the time of the transfer.

To facilitate trusted measurement, storage, and reporting of integrity metrics, the TP relies on three roots of trust. These are the Root of Trust for Measuring integrity metrics (RTM), and the roots of trust for storing and reporting integrity metrics. These roots of trust enable entities to trust the TP. The roots of trust are described further below.

6.6.1 TPM Identities

A TPM identity is used to attest to aspects of the TP. A TP can use this identity to prove that it is a valid TP.

One example of such an identity is called a TPM Identity Credential. This credential contains an identity label and a public key certificate, signed by a privacy Certification Authority (CA). The privacy CA attests that the identity label shown in the TPM Identity Credential belongs to a particular TP. The main role of TPM Identity Credentials is to protect user privacy by providing anonymity for the TP; for further details see, for example, [60].

Before a privacy CA signs a TPM Identity Credential, the TPM must first prove that the TP possesses certain properties. This proof is based on three credentials, in the form of certificates. These are the TPM endorsement credential, the platform credential and the conformance credential. These credentials have the following properties.

TPM Endorsement Credential. This credential attests that a Trusted Platform Module (TPM) conforms to the TCG specifications. The credential is signed by an entity such as the TPM manufacturer or a TPM conformance laboratory. The entity responsible for generating this credential is known as the Trusted Platform Module Endorsement (TPME) entity. By generating this credential, the TPME vouches that the TPM conforms to the TCG specifications. The credential contains the public key (PUBEK) of an endorsement key pair. This credential can be used by the privacy CA when generating a TPM Identity Credential to create an identity for the TP. The corresponding private key is called PRIVEK, and is kept in a shielded location in the TPM for confidentiality and integrity protection.

Platform Credential This credential attests that the platform as a whole correctly incorporates protected capabilities and shielded locations within the TPS. The entity that generates the Platform Credential, known as the Platform Entity (PE), may be the platform manufacturer or a conformance laboratory. The information found in this credential is used by the privacy CA when generating a TPM Identity Credential. Access to this credential is restricted for privacy reasons.

Conformance Credential. The Conformance Credential attests that the design of the TPS and its incorporation into the TP conforms to the TCG specifications. This may be generated by the platform manufacturer or a conformance

test laboratory. The information found in this credential is used by the privacy CA when generating a TPM Identity Credential. Access to this credential is restricted for privacy reasons.

6.6.2 TCG Measuring, Reporting and Storing Processes

As mentioned earlier, the RTM is the point of trust from which all integrity measurements are derived. The Core Root of Trust for Measurement (CRTM) is a component of the RTM, and is the point from which all trust in reporting measured information is derived. When the platform starts up, the CRTM makes an integrity calculation of the first component to be executed on the TP. This is reported to a Platform Configuration Register (PCR). The measured component then becomes the RTM, and is responsible for measuring the integrity of the next component to be executed. This measurement is then stored in a PCR and, again, the measured component is responsible for measuring the next value. This process then continues as the components of the TP execute.

An entity that challenges the TP in order to determine whether it is in a trusted state, receives in return PCR values together with corresponding validation data. Validation data contains values which should result when integrity measurements are made by a TP, and is signed by a trusted third party. For example, validation data may vouch for the integrity of particular software. Using this validation data, the challenger can verify the PCR values sent by a TP in response to a challenge. If the PCR values received from the TP are not the values expected by the challenger, based on the trusted validation data, then the TP may not be in a trusted state.

A TP may choose only to send PCR values in response to a challenge from a remote challenger if it first successfully authenticates this challenger. However, in general, validation data itself is not confidential information. It describes the values which should be produced if the platform is working correctly.

The PCRs are responsible for storing integrity values from the time that the platform was first booted up. We next describe the mechanisms used to compute and report on the values stored in the PCRs.

The *TPM_Extend* Operation This operation records new integrity measurement values for the PCRs at boot-up. The method used for this operation is described below.

When the trusted platform boots up, a *TCG_PCRVALUE* is initalised to zero. PCR_0, the first PCR value, is then set to 0 concatenated with a representation of the first event to be recorded. This representation is known as the *inDigest* value. The next *PCR* value, PCR_1, is the message digest of PCR_0 concatenated with the *inDigest* of the second event. This procedure continues. The TCG specification states that there must be at least 16 PCRs available for use to store these integrity digest values.

The *TPM_Quote* Operation This operation provides cryptographic reporting of PCR values. The values are signed using the private signing key corresponding to the public verification key found in the TPM Identity Credential. As mentioned earlier, a TP may have many TPM Identity Credentials, corresponding to its collection of identities.

6.6.3 Sealing Data

The TCG specification provides mechanisms used to dictate the platform state that must hold in order for encrypted data to be decrypted. We next discuss the use of this feature with regard to the release of personal data. This feature may also be used to protect cryptographic keys.

The mechanism relies on three objects, the `digestAtCreation`, the `digestAtRelease` and a set of PCR values. The PCR values used are those which correspond to the aspects of the platform which should be considered when data is to be released. This is not necessarily all the PCRs. The `digestAtCreation` is an integrity value calculated from the set of PCR numbers and their corresponding PCR values, when the sealed data item was created. The PCR numbers indicate particular PCRs. The `digestAtRelease` is an integrity value calculated from a set of PCR numbers and corresponding PCR values indicating the state of the platform that must hold in order for the sealed data to be released (i.e. to be made available in decrypted form). These three objects, together with the plaintext data, are all encrypted using the *TPM_Seal* function of the TCG specification. The sealed item is then stored on the

platform, together with the set of PCR numbers found in the `digestAtCreation`.

When a request is made to unseal this data, the TPM decrypts the sealed item. The data held within this sealed item is only released if, when using the listed PCRs, the recalculation of `digestAtRelease` corresponds to the values found in the sealed item.

TCG_Seal This operation is used to encrypt data. Addition information may be provided to ensure that this data is only released if a platform is in a pre-specified state.

TCG_Unseal This operation is used when the release of encrypted data is requested. The release of data using the *TCG_Unseal* operation may or may not depend on a pre-specified platform state.

6.7 Protecting Personal Information Using Trusted Computing

We now show how the privacy of personal information can be protected using the operations given in the TCG specifications.

6.7.1 Overview

Typically, when a subject transmits PI to a remote entity, it is then held in a database belonging to the receiving entity, in our case the SP. It would clearly be desirable that the SP's database stores PI in a manner which assures the subject that the PI is properly protected. Clearly, additional assurance is obtained if the user has guarantees of the platform state at the SP when its PI is used in the future. This ensures that the future use of PI depends on the platform being in a pre-defined trusted state.

This brings us to our first question. How does the subject know that the SP is in a trusted state? To answer this question we refer to the TCG specifications.

Assuming the SP is a TP, there are a host of mechanisms at our disposal. Using these mechanisms a subject can establish various properties of the SP platform. This includes identifying the software which has executed since the platform boot. The subject can then ensure that the software used is trustworthy.

This leads us to our next question: how does the subject know which software to trust? Knowledge of which software can be trusted could be established by a suitable test laboratory, and the properties of the software made available to the subject. This could include checks for the way in which personal data is managed by the software, and those security aspects of the software that affect the likelihood of unintended distribution of PI. The user can thus discover if the software used passes its criteria for trustworthiness. Of course the laboratory making the statements would itself need to be a trusted entity.

The process of ascertaining the software state on the target platform is initiated by a request for a proof of the state of the platform to which the PI is being sent. This request is in the form of a challenge. The platform then responds with a signed version of the challenge and PCR values, together with validation data. The private signing key used corresponds to the public key in the TPM identity certificate. The use of this key assures the user that the TP with which they are interacting has been deemed a valid TP by the TPM certificate provider. The inclusion of the challenge in the signed response prevents replay attacks. The validation data allows the user to recalculate the PCR values found in the signed response. The validation data may be provided by the testing laboratory, specifying the state that the software should be in if it is to be trusted. Trust in software may also be derived from the reputation of the software vendor. When the PCR values are recalculated, they can be compared to those sent by the SP. If they are the same, the user is assured of the target platform state and that of the software executing on it. The user can then determine whether or not to trust the software.

Given that the user now knows exactly which software is being used by the SP, the user can then decide whether or not to send its personal data to this entity. Additionally, the user can specify that the platform must be in a trusted state whenever the personal information is used in the future. This is guaranteed by sealing personal information, as discussed in section 6.6.3. The user could, for example, simply state

that their private information may only be used if this particular trusted software is executing on the platform.

6.7.2 Practicality

In this section we discuss some possible objections to our proposals. One of the main concerns with the TCG specifications is with the practicality of attestation. We discuss how this might affect the scheme that we have proposed here. We also look at schemes that attempt to address this issue.

The software that makes use of PI is likely to be different at every SP. If this is the case, then the practicability of the solution appears questionable, since it would seem to require every SP to have its software checked by a trusted third party laboratory. We make two observations about such a requirement. Firstly, if the SP is a major corporation, then this is probably not such an onerous task. Secondly, whilst the software that actually uses information derived from the PI to provide a service is likely to be SP-specific, the software that stores and processes the PI does not necessarily need to be. That is, one could imagine cases where the PI is handled by a trusted 'standard' application running on the SP server, the integrity and trustworthiness of which can be externally checked. This trusted application then only releases that part of the PI needed by the SP to provide its service, which might mean that the SP application itself does not need to be trusted.

The second main concern is how to deal with changes in software. Changes to software occur frequently, e.g. when software is upgraded, or a patch is added. When this happens, if the software is to remain verifiable, the SP must obtain new validation data for the new version of the software. Schemes that have been designed to address this issue use an approach that attests to the behaviour of the software rather than the binary code. Haldar, Chandra and Franz [65], use a virtual machine that attests to the behaviour of the software code. This virtual machine can attest to certain properties of the software when it executes in this environment, or it could analyse the software before it is executed. Sadeghi and Stüble [122] propose a scheme that uses a trusted attestation service. This also attests to certain properties of the software; however, instead of using a virtual machine, they propose that this

service is provided by an external entity. These properties, for example, may mean that the software complies with certain privacy regulations. The use of organisations for compliance testing is not a new concept. For example, the Wi-Fi Alliance attests that wireless devices comply with the 802.11 specifications.

The SP may not want to divulge information about the configuration of its software. This would provide requesters with information about the software that is installed on a platform and also the level of patching. This information could then be used for malicious purposes. Li, Shen and Zuo [90] propose the use of policies for trustworthiness to address this issue. These policies can be compared with the behaviour of software. This ensures that the trustworthiness of the software can be discovered without divulging any information about the software itself.

6.7.3 Using Trusting Computing with PI

In this section we consider how a subject acquires information about an SP platform. In particular, we discuss the TCG mechanisms which may be used by a subject to assess software found on the SP platform. That is, we describe in greater detail how the procedure outlined in section 6.7.1 may be implemented using TCG functionality.

6.7.3.1 Initiating a Request for PI

The first stage in our scenario involves an SP making a request to the subject for PI. This is usually the result of a request to the SP for the provision of a service to the subject. Before any information is provided to the SP the subject must establish certain properties of the SP's platform.

6.7.3.2 Verifying SP Platform State

When the subject receives a request for PI, it must ensure that the SP is in a suitable state to receive the PI. This will depend on the subject's judgement of the trustworthiness of the software found on the SP platform.

The TCG mechanisms allow the subject to ensure that a TCP is using trusted software. The subject can request information regarding the software found on a TP.

6.7.3.3 Sending PI to the TP

To check the software executing on a target SP platform, the subject must first transmit a random challenge to it. The SP may also require authentication of the subject.

Upon receiving the challenge, the SP performs the *TPM_Quote* operation described in section 6.6.2. This provides a signed version of the current PCR values for the SP platform, which is the result of the *TPM_Extend* operation. This operation is responsible for recording PCR values starting when the platform boots. The random value in the challenge is included within the scope of the signature generated by the *TPM_Quote* operation, preventing replay attacks.

Information regarding the software that has executed since platform boot is also sent. This is in the form of the *TCG_PCR_EVENT* data. Validation data, enabling the integrity of the software found on the platform to be verified, is also sent. The validation data is signed by the entity which created it. This may be a software vendor or a laboratory responsible for ensuring that software used for managing PI is secure and trustworthy. Using the information sent by the SP, the subject can recalculate the values calculated by the *TPM_Extend*, and compare them to the PCR values sent by the SP.

The subject must first validate the signed data and verify the TPM Identity Credential. If the signature and the credentials are valid, the subject can trust that the identity was created by the privacy CA. That is, at some time in the past, the SP has sent its credentials to the privacy CA, the privacy CA verified these credentials, and sent the SP a TPM Identity Credential. This is evidence that the privacy CA attests to the credentials presented by the SP. Of course, the subject must trust that the privacy CA performs its functions correctly.

The subject also recalculates the PCR values. If the result of this recalculation is as reported by the SP, the subject has assurances about the integrity of the SP platform state. Most importantly, the subject has assurances about the software used to manage PI and its integrity. If this software satisfies the conditions of the subject, it can then send its PI to the SP.

The subject may also request that PI sent to the SP is sealed. This means that the PI must be stored on the platform in encrypted form. Also contained within the encrypted data are objects which specify the conditions which the platform must satisfy in order for the data to be released. This ensures that, when the PI is accessed in the future by the SP, the SP platform is in a trusted state.

6.7.4 Constraints, LIPA and LI Tokens

As discussed in Chapter 4, constraints are statements which control the use, storage and distribution of LI. We next show how an LI software constraint can be included within the larger set of constraints. This special type of constraint can then be used to ensure that LI is only sent to LBS providers running trusted software to manage LI. By ensuring that an LBS provider uses specific software, the LI subject has assurance that the platform to which LI is sent will manage LI in a trustworthy manner.

We suppose that constraints are sent within the scope of an LI token, as discussed in Chapter 5. We further suppose that an LI token is encrypted in a manner such that only the LIPA is able to view the LI and the associated constraints. The LI token may then be distributed freely without any risk to subject privacy. When LI is required by an LBS provider, the LI token is sent to a LIPA. The LIPA decides whether or not the requesting LBS provider is authorised to have access to the LI contained in the token, based on the constraints contained in the LI token. If the LBS provider is a TP, the LIPA can check, using the mechanisms described above, that the LBS provider is using trustworthy software. The subject may also state in the constraints that LI must only be sent to LBS providers using trustworthy software. The presence of this software assures the subject that its LI is managed appropriately.

6.8 Other Approaches to Privacy Protection Using Trusted Computing

The use of trusted computing to address privacy issues has previously been proposed by a number of authors. Iliev and Smith [74] consider the use of trusted computing to enable Private Information Retrieval (PIR). PIR, first introduced in [26], aims to enable the acquisition of data, while keeping the data query private from the data holder. This may, for example, be used to ensure queries for medical records are kept private. This is potentially important for privacy reasons, since the mere existence of a query may indicate that a patient has an ailment.

Iliev and Smith's scheme proposes the use of a secure processor to randomly permute the data after it has been encrypted. The permutation is known only to the secure processor. When an entity wishes to acquire data, a request is made to the secure processor. The secure processor then fetches the data and sends it to the requesting entity over a secure channel. This ensures that the data holder knows nothing about the data which has been retrieved. When entities request further entries, the secure processor retrieves all previously retrieved entries together with the requested entry. If the request is the same as a previous request, all the previous entries and a random entry are sent to the secure processor. This ensures that the data holder does not know if requests made are the same as previous requests. Such systems could be used by consumers of LBSs to request LI from the holder of such information, without the holder knowing what was requested.

6.9 Analysis

This chapter does not discuss how users might find out what use an SP wishes to make of personal data. It may be possible for an SP to state that they wish to use PI for one particular purpose, whereas they may actually be using it for a different purpose. This issue of trust remains with any SP. One approach to tackling this issue is to transfer the risk to a different entity. This is usually the case with credit card transactions where credit card companies are usually willing to reimburse card holders when their card is used for fraudulent transactions. When considering PI,

there is no obvious entity to which this risk can be transferred. An alternative approach to raising awareness of mendacious SPs would be to to generate a list of those that cannot be trusted to handle PI in accordance with expressed policy. Such a list might, for example, be created by a third party. An entity may thus be able to establish with some degree of confidence that the service which the provider is claiming to provide is actually the service provided. Of course, such a list requires information from users that have previously had bad experiences. Such schemes are widely used and can provide a good basis for trust for SPs.

Another possible concern relates to how a user can establish whether or not the software used to manage personal information is trustworthy. Modern applications are constantly being modified and, in many cases, modifications are made by vendors to enhance the performance of software. In such a case, there is no malicious intent in the modification. However, any change in the software will, of course, change the integrity measurement. One possible solution to this problem may be to have independent testing laboratories which are responsible for ensuring that the software performs securely in the interests of the user. Of course this task may also be carried out by the original vendor.

A further problem which arises from this relates to differences in interests. The software vendor is selling its product to the service provider, so its primary concern is to keep its customer satisfied. How can the end user trust that the software vendor will act in its interests?

6.10 Conclusions

We have discussed how a TP may be used to gain assurance about the integrity of software found on a platform. This may be used by a subject to address concerns about the control of personal information. This can be done by verifying the integrity of software found on the platform to which personal information is being sent, and by ensuring that this software is trustworthy. The mechanisms described in this chapter rely on SP servers being TPs.

Chapter 7

Location-Based Services and the Usage Control Model

Contents

7.1 Introduction 113
7.2 LI Entities 115
7.3 Constraints 115
7.3.1 LI Gatherer Constraints 116
7.3.2 LI Distribution Constraints 116
7.3.3 LI Use Constraints 118
7.3.4 Difficulties in Implementing Constraints 118
7.4 An Introduction to UCON 119
7.5 Modeling LI Constraints in UCON$_{ABC}$ 121
7.5.1 Authorisation Restrictions 123
7.5.2 Conditional Restrictions 123
7.5.3 Model Definitions 124
7.5.4 Modelling LI Constraints in the UCON Model 126
7.6 Analysis, Future Work and Conclusions 129

In chapters 4 and 5 we introduced the notion of LI constraints, described how they may be used, and addressed some of the issues that may arise from their use. We also introduced the notion of a LI preference authority, a trusted third party whose role is to control access to LI and associated constraints. In this chapter we investigate the control of LI in the context of the UCON$_{ABC}$ usage control model. We look at how attributes of LI, and the entities in our architecture designed to control the dissemination and use of LI, fit within this model. We also look at the stages during the provision of an LI service at which LI control may take place. Some of the work in this chapter has been published in [55].

The entities involved in this chapter are defined according to the model developed in chapter 4. An overview of these entities is given in section 7.2. We describe how constraints can be used to control LI in section 7.3. Section 7.4 introduces the $UCON_{ABC}$ model. The use of the $UCON_{ABC}$ model to specify the use and management of LI constraints is described in section 7.5. Section 7.6 concludes the chapter by discussing the results and considering possible future work.

7.1 Introduction

There are various levels of private information; for example, a person may regard both a password and a telephone number as private information, although the password is likely to be regarded as "more" private. Another person may not regard a telephone number as private information at all. With this in mind, it is important to be able to control private information according to individual personal preferences. If a person does not want anyone to know their telephone number, they can simply not tell anyone. Of course, if the person wants to receive phone calls then this is not possible. In reality, personal information is distributed in a controlled manner, sometimes with attached constraints. For example, a person may only give their phone number to friends and family. If a person distributes their phone number to an unknown entity, then they may state restrictions on the way in which their personal information is to be managed. For example, when completing a form they may state that their personal information is not to be redistributed.

In this chapter we consider how constraints, i.e. statements of subject preferences regarding the handling of LI, can be modelled using a control framework recently proposed by Park and Sandhu [100, 101, 102, 123]. As discussed in Chapter 4, these constraints may be used to limit the intended use for the LI, the time at which LI may be obtained, or even the creation of the LI itself. Constraints are used in two main ways: they may be held by an entity (essentially as policy information) and constrain the way in which this entity handles LI, or they may be associated with an individual piece of LI and constrain the way that this piece of LI is handled.

We are by no means the first to investigate access control for LI. As discussed in Chapter 3, Leonhardt and Magee [89] generalise classical access control models [15,

87] and apply them to LI. They consider a variety of access control models, looking at using the notion of authorisation to control access to LI. For their purposes authorisation is based only on the identity of the requesting entity. Hengartner and Steenkiste [70] discuss controlling access to LI based on the location of the subject and the time at which the LI is requested. They also discuss implementation of their model using SPKI/SDSI certificates. In their model, control of LI is enforced in terms of its distribution, i.e. the entities to which LI distribution is permitted. Gathering of LI or the use of LI is not discussed.

In previous chapters we introduced the notion of LI constraints, described how they may be used, and addressed some of the issues that may arise from their use. We also introduced the notion of a LI preference authority, a trusted third party whose role is to control access to LI and the associated constraints. Traditional access control models simply cover authorisations and do not address other possible requirements for access control, such as conditions and prerequisites for access. These aspects are potentially useful in modelling access control for LI. For example, access to LI may be dependent on the time of day or the location of the requestor, and not just the identity of the entity requesting LI. This chapter looks at access control for LI using the UCON_{ABC} model [100, 101, 102, 123, 153], which extends traditional access control based on authorisations, and provides the controls that we desire.

Park and Sandhu have developed a general access control model called the usage control $(\text{UCON})_{ABC}$ model. This model is designed to address issues such as Digital Rights Management (DRM), code authorisation, and the control of personal data. UCON extends the traditional notion of authorisation by introducing *obligations* and *conditions*. An obligation requires a subject to perform a task in order to exercise a right on an object. For example, a person may have to sign a waiver before being allowed to participate in a parachute jump. A condition requires a predefined state to be true before the subject can exercise a right on an object. For example, the temperature must be below 15°C before the heating is permitted to be turned on. Using a combination of authorisation and the Park-Sandhu conditions, in this chapter we describe the use and management of LI constraints in the context of this general UCON_{ABC} model. The aim of this chapter is to use our notion of constraints to test the UCON_{ABC} model.

7.2 LI Entities

Figure 7.1 shows the different stages in the provision of an LBS. Constraints may be placed on the use and management of LI at each of these stages. At the LI gathering stage, an LI subject may constrain the times when gathering may take place. Constraints may also be placed on the distribution of LI. For example, only certain entities may be permitted to possess LI. The entity responsible for distributing LI at the initial stage may be the LI gatherer; however, further distribution may be performed by entities such as the UD of the LI subject, a network entity, the LBS provider, or an LBS directory. Once an entity receives LI, it may wish to use it. Of course, to use the LI, the entity must have the necessary usage rights, as specified by the associated constraints.

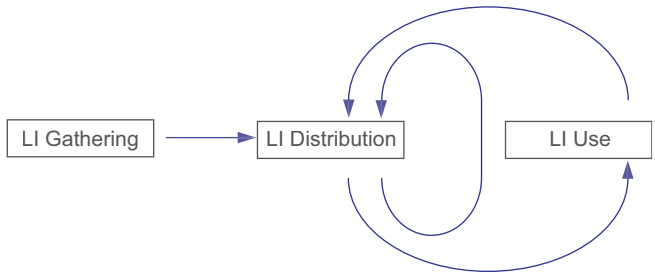

Figure 7.1: The stages at which control can be applied to LI

7.3 Constraints

In Chapter 4 we introduced the use of constraints, which are simple statements used to control the dissemination and use of LI. In most cases we would envisage constraints being defined by the LI subject. As discussed in Chapter 4, entities such as a regulatory authority may have requirements that contradict the constraints set by the LI gatherer. In such cases the requirements of the regulator must be considered first. For example, in some countries [68] the provision of an emergency service may take precedence over any constraints which may be in place.

Constraints may be held in a central server or, alternatively, transmitted securely with LI. These issues are discussed further in Chapter 5. Constraints can be used

to control the use and distribution of LI at three stages in the provision of an LBS: when LI is first generated by the LI gatherer; when LI is distributed by an entity; and when LI is used, typically by an LBS provider. The types of constraint used at each of the stages are now described.

7.3.1 LI Gatherer Constraints

LI gatherer constraints limit the gathering of LI. We give three classes of LI gatherer constraints, although such constraints are not limited to these types. Note also that, if the gatherer constraint is particularly restrictive, it may prevent the gathering of any LI.

- Gathering time constraints limit the time during which LI can be gathered. For example, a user may only wish LI to be gathered by his employer during working hours.
- Gathering accuracy constraints may limit the accuracy of LI gathered. For example, the LI gatherer may only be permitted to gather LI indicating the location of the LI subject to an accuracy of the nearest kilometre.
- Constraining the gathering of LI based on the location of the LI subject would be difficult, because only once LI is gathered can the LI subject's location be assessed. Constraints based on subject location are considered in section 7.3.2.

7.3.2 LI Distribution Constraints

LI distribution constraints limit the entities to which LI is distributed. These constraints can be further subdivided into distribution time constraints, distribution entity constraints, distribution usage constraints, distribution location constraints and distribution validity constraints. Such constraints may either be bound to individual LI records, i.e. pieces of LI, in which case they control the distribution of that record only, or they may be used to control the distribution of all the LI relating to a particular subject. In the latter case such a constraint would be lodged with one or more entities responsible for disseminating LI.

Distribution time constraints can be used to express limitations on the times at which LI may be distributed. For example, an LI subject may not want LI distributed during work hours.

Distribution entity constraints control the entities to which LI is sent. Such constraints can be specified inclusively or exclusively; inclusive constraints specify the entities who are permitted to possess LI, whereas exclusive constraints specify the entities who are not permitted to possess LI.

Distribution usage constraints prevent the distribution of LI to entities which may use it for a purpose not desired by the LI subject. For example, an LI subject may not want its LI distributed to entities intending to use it to advertise to them. For example, when the LI subject is walking past a certain shop, the shop owner may wish to advertise a sale to the LI subject, but the LI subject may not wish to receive such advertisements.

LI may be used to provide the LI subject with a service or with location details, or provide a service or location details to a separate entity. For example, location-based messaging may be used to inform an LI subject regarding any buddies who may be located nearby, where these buddies may be members of a list maintained by the LI subject. This is similar to internet messaging with the added location property. An LI subject may not want buddies to know his/her location at a given time, a restriction which can be supported by the use of appropriate LI constraints.

Distribution validity constraints control the distribution of LI based on an assigned validity period. A constraint of this type will either specify how long LI is valid after generation, or will specify its time of expiry; once the time period has expired the LI must not be distributed. This type of constraint would typically be bound to an individual LI record, since each piece of LI will have an individual expiry time.

Finally note that there are three possible effects of a distribution constraint with respect to the distribution of a single LI record.

- First, the LI record may be permitted to be distributed unchanged.
- Second, distribution of the LI record may be completely prevented.

- Third, and most interestingly, the constraint may only permit the distribution of the LI record in a modified form, e.g. after it has been made less accurate (in time or space), or less specific, e.g. referring only to the LI subject by a group of which it is a member, or a role which it holds.

7.3.3 LI Use Constraints

LI use constraints are used to control the use of LI, e.g. by an LBS provider. These restrictions include usage validity constraints, usage location constraints, usage time constraints, usage purpose constraints and possession constraints.

Usage validity constraints control the use of LI based on its validity period, as specified in a constraint bound to an individual LI record. For example, some LBS providers may not be permitted to use LI if it is due to expire within a certain period of time.

Usage location constraints control use of LI based on the location to which the LI refers. For example, an LI subject may not want to be locatable when they are at home.

Usage time restrictions prevent LI from being used at certain times. For example, this may be used to prevent LI from being used during non-working hours.

Usage purpose constraints limit the use of LI to particular purposes (or deny its use for specified purposes). For example, an LI subject may not wish his or her LI to be used for advertising purposes.

Possession constraints prevent LI from being used by an entity not permitted to do so.

7.3.4 Difficulties in Implementing Constraints

The LI gatherer must be a trusted entity. When LI is gathered on a UD, the LI subject can apply constraints to its LI directly. LI may also be gathered by a network

entity; in this case the LI subject must trust it to act according to the constraints.

Difficulties arise when LI is distributed or used by other entities. Once an entity possesses LI, retaining control of it is difficult, as they must be trusted to apply any controls. Trusted computing platforms could potentially force entities to obey constraints. The use of trusted computing platforms in relation to LI and privacy is discussed further in Chapter 6. Enforcing constraints is not discussed further in this chapter.

7.4 An Introduction to UCON

The $UCON_{ABC}$ [100, 101, 102, 123] model contains eight components, namely subjects, objects, rights, subject attributes, object attributes, authorisations, obligations and conditions. These components form the basis of a generic model to control access to resources.

In this model, the subject, which possesses subject attributes, is an entity which wishes to gain access to an object, which itself has associated object attributes. Subject attributes and object attributes are the properties of the subject and object respectively. Access is based on the rights which a subject may exercise depending on the subject attributes and object attributes. A right could, for example, be one of read, write or execute. The precise nature of the rights will vary depending on the object and the context of access to the object. For example, in the context of LI, a right might be one of gather, receive or use. Authorisations, obligations and conditions are used to determine whether a subject may exercise a right upon an object. Authorisations determine if a subject is permitted to exercise a right on an object. An obligation specifies a task that a subject must execute in order for it to exercise a right on an object. Finally, a condition specifies a state that must be satisfied in order for the subject to exercise a right on an object.

As defined above, the access decision-making process to exercise a right will be based on a comparison of the subject attributes with the object attributes, bearing in mind whether or not the authorisations, obligations and conditions have been met. We look at the possible use of this model to control the use of LI.

The authorisation property asks the question, is the subject allowed this right to this object? To answer this question a decision-making process must establish whether the subject has the correct subject attributes to access the object, given its specific object attributes. An authorisation may be a pre-authorisation, *pre*A, or an on-going authorisation, *on*A. If an authorisation is *pre*A, the authorisation process takes place before the right is exercised on the object. If an authorisation is *on*A, authorisations take place continuously. That is, before and during the exercise of the right.

Obligations ask the question, has process A taken place so that a right may be exercised? For example, a subject may have to sign a contract before being allowed a right to some information. Obligations may also take the form of a pre-obligation, *pre*B, or an on-going obligation, *on*B. We do not discuss obligations when describing use of the $UCON_{ABC}$ model with LI.

Finally, a condition property asks the question: is condition B true so that a subject may have the requested right to this object? Again, a condition make take the form of a precondition, *pre*C, or an on-going condition, *on*C.

An important property of the UCON$_{ABC}$ model is that it allows mutability. This means that attributes may be changed before, during or after a right has been exercised. Four levels of mutability are described. These are: immutable; pre-update; ongoing-update; and post-update. If an attribute is immutable, it may not be changed. Pre-update means that the attribute can be changed before a right to an object is exercised. Ongoing-update means that an attribute may be constantly changing. Post update means that the attribute can be changed after access to the object has been gained.

Park and Sandhu [100, 101, 102, 123] identify 16 key instances of the UCON$_{ABC}$ model; these are referred to as the 16 'core models'. Each of these instances corresponds to a particular choice of mutability properties for the three types of access control attribute; the core models are represented by the Ys in Table 7.1. The first column represents the case where none of the attributes are mutable, regardless of whether they are *pre* or *on*. The second column represents the case where the attribute is updated before a right is given. Observe that *pre*C and *on*C are the only

Table 7.1: Instances of UCON$_{ABC}$

	0 (immutable)	1 (pre-update)	2 (ongoing-update)	3 (post-update)
preA	Y	Y	N	Y
onA	Y	Y	Y	Y
preB	Y	Y	N	Y
onB	Y	Y	Y	Y
preC	Y	N	N	N
onC	Y	N	N	N

attribute types which cannot be mutable before a right is given. Conditions are not generally mutable. For example, if an object may only be accessed at a certain time, the time attribute is not likely to change after the object has been used. The third column in table 7.1, ongoing-update, refers to the case when the attributes may be mutable before, during or after a right is exercised. The fourth column, post-update, represents the case where the attribute is mutable after a right has been exercised. This is possible for all cases except where a condition is used for the decision. As with the pre-update column, conditions are not generally mutable.

7.5 Modeling LI Constraints in UCON$_{ABC}$

In this section we apply the UCON model to the LI constraints architecture defined in Chapters 4 and 5. That is, we will describe the use of LI constraints to control LI in a formal way. There are two reasons for doing this. First, it gives a formal basis for the LI constraints model. Second, it provides an interesting use case for the application of the Park-Sandhu model.

As mentioned in Section 7.3, LI access can be controlled at three stages. That is, when LI is first generated, when LI is requested by an entity, and when LI is used. We categorise the constraints into authorisations and conditions, and describe the constraints used at each stage. We do not map any of our constraint types into obligations, since the notion of an obligation does not appear to arise naturally when considering the control of the use and management of LI. In Table 7.2 the various constraint types identified in Section 7.3 are mapped to Authorisations and Conditions in the UCON model. Note that we do not map gathering accuracy

Table 7.2: Mapping LI constraint types to UCON attribute types

	Authorisation	Condition
LI gathering		*Gathering time*
LI distribution	*Distribution entity*	*Distribution validity*
	Distribution usage	*Distribution time*
LI use	*Usage purpose*	*Usage validity*
	Possession	*Usage time*
		Usage location

constraints, since these appear hard to model using UCON$_{ABC}$.

The mapping of the LI constraint model outlined above to the UCON model has the following properties.

- The UCON$_{ABC}$ model uses the term 'subject' to denote the entity which wishes regarding the exercising of a right on an object. In our application of the UCON model, the UCON subject corresponds to an LI handling entity, i.e. an entity that will gather, receive or use LI. The subject attributes, used in access control decisions, correspond to LI constraints applying to the LI handling entity. It is important to note that the subject in the UCON model is very different from the LI subject; in fact the LI subject has a close relationship to the object in the UCON model (see below).
- In the LI application of the UCON model, the objects correspond to individual pieces of LI. The object attributes then map to constraints bound to individual pieces of LI.
- As specified above, a right indicates an action which a subject may invoke on an object, e.g. read, write and execute. The rights which we use in our application of the UCON model are gather, receive and use. If a subject wishes to exercise a right on an object, the authorisation, obligation and condition properties must be satisfied. These properties can be found in the subject attributes and the object attributes. Of course, it is not always necessary to use all these properties together.

7.5.1 Authorisation Restrictions

Authorisations will be verified before LI is distributed, and before LI is used. The LI constraints which involve authorisations are distribution entity constraints, distribution usage constraints, usage purpose constraints, and possession constraints. Before LI is distributed, the distributing entity compares the identity of the entity to which LI is being sent with the distribution entity constraint, and the intended use of the LI by the entity to which the LI is being sent with the distribution usage constraint. This establishes whether the receiving entity is eligible to possess this LI. Both the distribution entity constraint and the distribution usage constraint are pre-authorisations, i.e. the authorisation takes place only before the distribution of LI. Of course any authorisation which takes place after the LI has been distributed is pointless, because at this point the entity already possesses the LI.

Before the LI is used, the entity must itself establish if it is permitted to use the LI by examining the usage purpose constraint, and if it is permitted to possess the LI by considering the possession constraint. An entity may be permitted to possess LI but not be permitted to use it if, for example, its role is only to distribute the LI. Usage purpose constraints and possession constraints are also pre-authorisations. The authorisations described are immutable, i.e. they do not change as a result of exercising a right.

7.5.2 Conditional Restrictions

Conditional restrictions refer to constraints which take into account the circumstance in which a right is being exercised when making a control decision. Conditional restrictions are used when LI is gathered, distributed and used. The conditional restrictions listed in Table 7.1 are gathering time constraints, distribution validity constraints, distribution time constraints, usage validity constraints, usage time constraints and usage location constraints.

When LI is gathered, the time when LI gathering takes place must be consistent with the times at which LI gathering is permitted by the LI subject, where the permitted times are recorded in the gathering time constraints. This is a pre-condition, i.e.

this is decided before the LI is gathered.

When LI is distributed, the constraints that need to be considered are the distribution validity constraints and the distribution time constraints. These conditions are also pre-conditions. The distribution validity constraint specifies the duration for which the LI may remain in existence. For example, an LI subject may only want LI to exist for an hour after it has been gathered. If the LI is no longer valid then it should be deleted by the entity distributing it. The distribution time constraint specifies the times at which LI may be distributed.

When LI is used, the constraints that must be considered are the usage validity constraint, usage time constraint and usage location constraint. The usage validity constraint specifies the validity status of the LI that must hold in order for it to be used. The usage time constraint specifies the time at which LI may be used. The usage location constraint specifies the set of locations which, if the LI information is a member of the set, constrain its use. For example, if an LI subject is at home, then it may not want its LI used at all. Usage time constraints and usage validity constraints are on-going conditions, i.e. these constraints have to be evaluated continuously as the LI is used. This is because the current time will change during LI use, and a constraint which allows LI use at one point in time may deny it at another.

The usage location constraint is not an on-going condition. Although the LI subject may change its location while the LI is being used, usage location constraints refer to the location information in the LI itself, rather than the actual current location of the LI subject. When new LI is sent to the using entity, access to the LI may then be denied if the usage location constraint reflects this.

7.5.3 Model Definitions

This section gives the definitions used to describe LI constraints. We subsequently model the use of LI constraints using these UCON$_{ABC}$ model concepts.

General definitions

S, O, R, ATT(S), ATT(O), preA, preC and onC

These represent subjects, objects, rights, subject attributes, object attributes, pre-authorisations, pre-conditions and on-going conditions, respectively.

Definitions for *pre*A (pre-authorisations)

allowed(s,o,r) $\Rightarrow$ *pre*A(ATT(S), ATT(O),r)

The left-hand side indicates that a subject s is permitted right r to object o. This implies the right-hand side, where *pre*A uses ATT(S), ATT(O) and r to establish a usage decision.

Definitions for *pre*C (pre-conditions)

preCON This is the set of pre-condition elements.

getPreCON : S $\times$ O $\times$ R $\rightarrow 2^{preCON}$ This is a function which monitors selected pre-condition elements.

preCONChecked : preCON $\rightarrow$ {true, false} This function checks that pre-conditions have been performed without any error.

*pre*C(s,o,r) = $\wedge_{\ preCON;\ \epsilon getPreCON(s,o,r)}$ preConChecked(preCon$_i$)

This function uses the preCONChecked function to ensure that the pre-conditions relevant to this right are checked.

allowed(s,o,r) $\Rightarrow$ *pre*C(s,o,r) A right is permitted or denied depending on the result of the *pre*C function.

Definitions for *on*C (on-going conditions)

onCON This is the set of on-going condition elements.

getOnCON : S $\times$ O $\times$ R $\rightarrow 2^{onCON}$ This is a function which monitors selected on-going condition elements.

onCONChecked : onCON $\rightarrow$ {true, false} This function checks that on-going conditions have been performed without any error.

*on*C(s,o,r) = $\wedge_{\ onCON;\ \epsilon getOnCON(s,o,r)}$ onConChecked(onCon$_i$)

This function uses the onCONChecked function to ensure that all the on-going conditions relevant to this right are true.

allowed(s,o,r) $\Rightarrow$ true A right is permitted as long as the above function is true.

stopped(s,o,r) $\Leftarrow \neg$*on*C(s,o,r) The stopped procedure is performed when the requirements are no longer satisfied.

7.5.4 Modelling LI Constraints in the UCON Model

This section describes the use of LI constraints in the UCON$_{ABC}$ model. We do not describe constraints where a right may be for LI of reduced accuracy.

7.5.4.1 LI gathering model

When LI is being gathered, we consider only gathering time constraints. We describe this as a pre-condition in the context of the UCON model. This means that the gathering time constraint must be evaluated before LI is gathered. Additionally, this condition is immutable. From Table 7.1 we can see that such a scenario is covered by the UCON$_{ABC}$ model. In fact this scenario is shown by the following convention: UCON$_{preA_0}$ with the 0 representing the immutable column and $PreA$ indicating the pre-authorisation. A specification indicating whether or not an LI gatherer is permitted to gather LI is given directly below. This should be evaluated before an LI gatherer gathers LI.

LI gatherer gathers LI

Preliminary descriptions:
Erole : Set of roles for entities
T : Set of times
currentT : Current time
gathtimes : O $\rightarrow$ T Times at which LI may be gathered
entityrole : S $\rightarrow 2^{Erole}$
ATT(S) = {entityrole}
ATT(O) = {gathtimes}

Main description:
preCON {currentT ϵ gathtimes(o)}
getPreCON(s,o,r) =
{currentT ϵ gathtimes(o) if entityrole(s) = LIGatherer}
allowed(s,o,gather) $\Rightarrow$ preCONChecked(getPreCON(s,o,gather))

7.5.4.2 LI Distribution Model

When LI is distributed, the pre-authorisations, i.e. the distribution entity constraints and the distribution usage constraints, are considered. Pre-conditions, i.e. distribution time constraints and distribution validity constraints, are also considered. Both the pre-authorisations and the pre-conditions are immutable. Both of these possibilities exist in the UCON model, as can be seen from Table 7.1. This is given by the following convention: UCON$_{preA_0preC_0}$. A specification indicating whether an entity is permitted to receive LI is given directly below. This should be evaluated before LI is sent.

Entity may receive LI

Preliminary descriptions:
E : Set of entities Erole : Set of roles for entities
T : Set of time
currentT : Current time
U : Set of uses for LI
entityrole : S $\rightarrow 2^{Erole}$
entityname : S $\rightarrow 2^{E}$ Name of requesting entity
entityuses : S $\rightarrow$ U Entities to which LI may be sent
disttimes : O $\rightarrow$ T Times at which LI may be distributed
validity : O $\rightarrow$ T Time when LI becomes invalid
canbesent : O $\rightarrow$ E Entities to which LI may be sent
canbeusedfor : O $\rightarrow$ U LI uses
ATT(S) = {entityrole, entityname, entityuses}
ATT(O) = {disttimes, validity, canbesent, canbeusedfor}
preCON {currentT ϵ disttimes(o), currentT $\leq$ validity(o)}
getPreCON(s,o,receive) = {
currentT ϵ disttimes(o) $\bigcap$ currentT $\leq$ validity(o)
if entityrole(s) = LIEntity}

Main description:
allowed(s,o,receive) $\Rightarrow$ preCONChecked(getPreCON(s,o,receive)), $\exists$ o'| can-

beusedfor(o') ϵ entityuses(s), entityname(s) ϵ canbesent(s)

7.5.4.3 LI Use Model

When LI is used, pre-conditions, i.e. usage validity, usage time, and usage location are considered. These pre-conditions are also considered as on-going conditions. This means that these restrictions will also apply during LI use. Pre-authorisations, i.e. usage purpose and possession are also considered. The pre-conditions, on-going conditions, and pre-authorisation are all immutable. This is given by the following convention: UCON$_{preA_0 preC_0 onC_0}$. A specification indicating whether an entity is permitted to use LI is given directly below. This should be evaluated before LI is used.

Entity may use LI

Preliminary descriptions:
E : Set of entities
Erole : Set of roles for entities
T : Set of time
currentT : Current time
U : Set of uses for LI
L : Set of locations entityrole : S $\rightarrow 2^{Erole}$
entityname : S $\rightarrow 2^{E}$ Name of requesting entity
entityuses : S $\rightarrow$ U Entities to which LI may be sent
usetimes : O $\rightarrow$ T Times at which LI may be distributed
validity : O $\rightarrow$ T Time when LI becomes invalid
canpossess : O $\rightarrow$ E Entities which may possess LI
canbeusedfor : O $\rightarrow$ U LI uses
usagearea : O $\rightarrow 2^{E}$ locations where LI may be used
li : O $\rightarrow$ L LI
ATT(S) = {entityrole, entityname, entityuses}
ATT(O) = {usetimes, validity, canpossess, canbeusedfor}
preCON {currentT ϵ usetimes(o), currentT $\leq$ validity(o),
*on*Con {currentT ϵ usetimes(o), currentT $\leq$ validity(o) }

li(o) ϵ usagearea(o) }
getPreCON(s,o,use) = {
currentT ϵ usetimes(o) $\bigcap$ currentT $\leq$ validity(o)
if entityrole(s) = LIEntity}
getOnCON(s,o,use) = {
currentT ϵ usetimes(o) $\bigcap$ currentT $\leq$ validity(o)
li(o) ϵ usagearea(o) if entityrole(s) = LIEntity}

Main description:
allowed(s,o,use) $\Rightarrow$ preCONChecked(getPreCON(s,o,use)), $\exists$ o'| canbeused-for(o') ϵ entityuses(s), entityname(s) ϵ canpossess(s)
stopped(s,o,use) $\Rightarrow$ $\neg$onCONChecked(getOnCON(s,o,use))

At the LI gatherer stage, the subject is the entity wishing to generate LI. The object in this case is a service and not actual data.

7.6 Analysis, Future Work and Conclusions

In this chapter, we have shown how some LI constraints can be described using the UCON$_{ABC}$ model. As part of the process of using this model, it has become clear that LI may be controlled at three different stages: when LI is gathered, distributed and used. We have also described some of the more important classes of constraints which may be used to control the management and use of LI.

Unfortunately, the UCON$_{ABC}$ model did not allow us to express controls when access to LI is denied. The two main examples of this occur when a right to LI is restricted based on its validity, and when it is necessary to reduce the accuracy of LI.

When a right to LI is denied based on its validity, a restriction should be imposed preventing any further attempt to use it. Although access to the object has been represented in the model, actions to be taken as a result of the denial of a right are not discussed. Further enhancement to the UCON$_{ABC}$ model should encompass

actions as a result of denial of a right. For example, in the case of the validity example, when LI is invalid it should be deleted, preventing its further use.

We also identified the possibility that LI of a reduced accuracy rather than the original LI may need to be provided in certain cases. Such modifications to the object being accessed are somewhat outside of the scope of the UCON_{ABC} model. The UCON_{ABC} model does, however, describe derivative objects which may be a means of describing this case. A derivative object is created as a consequence of exercising a right on an object. For example, a log file might be such a derivative object. In our case, if the LI of reduced accuracy was created when a right is denied, then this could be an alternative LI object for a subject. Of course the LI constraints would also have to apply to the alternative LI. Such a feature may also be useful in other scenarios. For example, when researchers request medical data, it is not always necessary to obtain data with identifying information (or access to only part of the data may be allowed). Access should then be granted only to the data without identifying information (or only to the allowed parts of the record).

Further work in this area might include devising a suitable language to represent LI constraints.

Chapter 8

Generating LI Using Ad Hoc Networks

Contents

8.1 Introduction **132**
8.2 Terminology **134**
8.3 Ad Hoc Networks **135**
8.4 The Location Discovery Service **136**
8.5 Requirements and Architecture **137**
8.5.1 Requirements 137
8.5.2 Infrastructure Based Tracking 138
8.5.3 Ad Hoc Tracking 139
8.5.4 Other Scenarios 140
8.6 Location Technology Overview **140**
8.6.1 The GPS Method 140
8.6.2 The Smart Antenna Method 141
8.6.3 DOA for Omnidirectional Antennas 142
8.7 Security Requirements **142**
8.8 Security Solutions **144**
8.9 Related Work **146**
8.10 Future Work **150**
8.11 Summary and Conclusions **151**

This chapter presents a generic service which allows an LI gatherer to discover the location of other devices in an ad hoc network. The service has advantages in a variety of scenarios, since it does not rely on location infrastructures such as GPS satellites or GSM cellular base stations. An outline of the technology that will be needed to realise the service is given, along with a look at the fundamental security issues which surround the use of this service.

We begin this chapter with an introduction in section 8.1. Section 8.2 introduces the terminology used in this chapter. A description of ad hoc networks and a proposal for using them as part of a locating service are given in sections 8.3 and 8.4. The requirements for such a service are described in section 8.5. The location technology which may be involved in providing the location service is described in section 8.6. Security requirements and solutions are discussed in sections 8.7 and 8.8. Related work and future directions are discussed in sections 8.9 and 8.10. Finally, we conclude this chapter in section 8.11.

The research described in this chapter is joint work with Po Wah Yau.

8.1 Introduction

The emergence of wireless technology has provided a catalyst for industry and academia to develop numerous new applications and services. How the underlying wireless technology works dictates what services can be provided. An example of this is the emergence of ad hoc networks as a communications medium. For our purposes, ad hoc networks are a permanent or temporary collection of nodes that can communicate with each other. The distinguishing properties of such networks are that there is no pre-existing infrastructure, there is no central entity to provide network administration services, and end-to-end communication may require information to be routed via several nodes[8.1]. A detailed discussion of ad hoc networks and their possible applications can be found, for example, in [107].

Ad hoc networks are potentially very useful in certain scenarios, such as emergency response networks, where a dynamic set of entities, such as police, fire services, paramedics or other agencies, need to intercommunicate in an environment where no communications infrastructure exists, either because there was none to start with or because it has been destroyed by a disaster. In this chapter we present a service that can be provided in an ad hoc network environment that enables the location of an object to be determined by appropriately authorised users. This is achieved

[8.1] This is why ad hoc networks are sometimes referred to as multi-hop networks, where a hop is a direct link between two nodes. If wireless communications are being used, then two nodes are within one hop of each other if they lie in each other's transmission range.

using ad hoc network routing principles, so that there is no need for an expensive communications infrastructure.

In previous chapters, LI is provided to an LBS provider who uses this information for the provision of a service. In this chapter, we investigate a user in an ad hoc network that wishes to locate a device. In scenarios used in previous chapters, the LI subject has been human user. In the examples described below, the LI subject refers to devices that are being located.

The first scenario involves the use of the service to locate a vehicle. One case where such a service would be useful is where a driver walks into a car park but forgets where his car is parked. The driver's mobile phone can form an ad hoc network with all the cars in the car park, including the driver's car and other ad hoc capable devices. On request, the driver's phone can broadcast the car's identifier[8.2]. The network can then respond to the request, providing the user with details on where the car is parked. Another case where such a service would be useful is if the driver's car is stolen. Depending on the scale and pervasiveness of ad hoc network nodes, the location service could be used to track the vehicle, giving a potentially valuable tool for the police service. This application could thus provide a low cost alternative to expensive tracking devices which use the GPS satellite system [45, 75].

A second scenario involves locating items of stock in a warehouse. We suppose that the stock items contain devices capable of forming an ad hoc network. When a warehouse worker wishes to locate an item in the warehouse, the stock items create an ad hoc network which is used to indicate the location of the desired item to the warehouse worker. Such a scenario could vastly reduce the time and cost of stock-taking. Instead of itemising the goods found in the warehouse by going through the laborious process of checking each individual item, the process could be automated by checking the nodes of the ad hoc network.

A third application is in military scenarios, appropriately given that research in the use of ad hoc networks was originally driven by such scenarios. The ability to accurately locate military devices and personnel has obvious advantages in battlefield scenarios.

[8.2] Or some other information identifying the car whose location is being sought.

Yet another set of applications is provided by the 'active office' environment [67, 149]. Here, users or even an automated telephone system can locate where colleagues are located within an 'active' building, e.g. to route telephone calls. Alternatively, a user's PC work environment might be automatically transferred to a display adjacent to their location.

Possible ways in which the service can be provided are outlined in this chapter, along with a review of what underlying location discovery technologies might be appropriate to support the service. The location discovery service has some potentially very important security and privacy requirements which are also discussed, along with initial thoughts on how these requirements can be met.

8.2 Terminology

The following terms are used in this chapter, but may be used differently elsewhere. A *node* is a device which has a network interface that is participating in the ad hoc network's routing service. It may or may not be mobile, and may also be part of another network. It is important to realise that a node can actually be a large network, or it could just be a single mobile device such as a mobile phone. An *LI requestor* is a node which wishes to discover the location of other nodes, known as *LI subjects*.

A node is a *neighbour node* of another node if it is only one hop away and within direct transmission range. If the destination node is not a neighbour node of the originator node, the data packet will have to traverse a multi-hop route consisting of *intermediate nodes*. In a specific scenario, the *sending node* is the last node to have forwarded the data packet.

There are two types of location discovery. The first is *actual LI*, where an LI requestor learns the exact geographical location of an LI subject, to a certain degree of accuracy. This is also discussed in Chapter 1. The second is *relative LI*, where an LI gatherer will discover the location of the LI subject relative to its own location, e.g. in terms of which direction the LI subject is located.

8.3 Ad Hoc Networks

The motivation for using ad hoc networks for this location discovery application is that ad hoc networks have the potential to be deployed anywhere, leading to true pervasive computing. They are thus not subject to environmental limitations which may prevent other technologies from working. Also, the multi-hop nature of ad hoc networks means that each device does not need sophisticated and potentially power hungry wireless communications facilities to be able to exchange information with the whole network. We assume the existence of an ad hoc network independent of any infrastructure, although there may, of course, be limited infrastructure available.

This service makes use of ad hoc routing protocols to disseminate information. We will describe an application protocol to implement the service that runs over an ad hoc network. This creates various requirements on the underlying network architecture. These are discussed further in section 8.5.

There are two main types of ad hoc network routing protocol, namely pro-active and reactive protocols. Within these categories, individual schemes use a variety of techniques to find and maintain routes. Most routing protocols are table-driven, where information is processed and stored in routing tables.

Reactive protocol operation is typically divided into a route discovery cycle and route maintenance. A node initiates route discovery when it needs to send a data packet to a destination for which a route is not known. This typically involves broadcasting some form of route request message, where an intermediate node or the destination node itself can provide the originator node with a reply containing the route to the destination. Route maintenance is required, as there are no periodic route update messages. Instead, when a link break is detected between two nodes, one or both of these nodes are responsible for propagating error information about the broken link to all affected parties. Examples of reactive routing schemes include the Ad hoc On-Demand Distance Vector (AODV) protocol [106]; Dynamic Source Routing (DSR) [81], which uses 'source routes'; and Location Aided Routing (LAR) [144], which uses geographical coordinates to increase the efficiency of routing.

Pro-active protocols use periodic topology updates to disseminate route information

throughout the network, but try to minimise the information being sent in order to save bandwidth. Various techniques are used to achieve this, as exemplified by the Optimised Link State Routing (OLSR) [30] and Topology Broadcast Reverse Path Forwarding (TBRPF) [16] protocols.

8.4 The Location Discovery Service

We now give an overview of the service and introduce some terminology. An outline is then given of possible technologies that may be used to provide the service.

We suppose that the LI requestor is a user that has one or more wireless-enabled devices, perhaps as part of a Personal Distributed Environment (PDE) [41]. When the user wishes to locate a device beyond the radio range of its own device, it can do so using one of the devices in an ad hoc network.

The LI requestor may or may not be currently operating in the ad hoc network being used to provide the location service. If not, the user must first perform whatever operation is required to make the device join this ad hoc network, including providing any necessary authentication information. Once this has been achieved, the user will need to specify the identity of the device to be located.

The location discovery service is provided using a special pair of messages sent through the ad hoc network. The LI gatherer broadcasts the identifier of the LI subject throughout the ad hoc network using a *TrackingRequest* message[8.3]. When the LI subject receives the *TrackingRequest*, it unicasts a *DirectionReply* to the LI requestor. This *DirectionReply* is forwarded back to the LI requestor via intermediate nodes. When the LI requestor receives the *DirectionReply* it uses LI contained within the message to determine the direction and distance of the LI subject. The contents and format of the LI will depend on the underlying technology being used, and this is discussed further in section 8.6.

As the nodes may be mobile, the service could be periodically re-run, so that the LI subject periodically sends a *DirectionReply*. To save power, the LI subject could

[8.3] This is equivalent to a Route Request message in a reactive ad hoc routing protocol.

even be instructed to sleep, checking less incoming messages. It could be instructed to wake when it expects to be located by the LI requestor.

Possible advantages offered by this application include that it allows smaller devices with restricted battery power to participate, and not every device needs location aware hardware such as a GPS receiver. The service is designed to cope without an infrastructure, but is capable of taking advantage of an infrastructure should it be available. Section 8.9 gives a brief description of how similar services are offered by other technologies.

The success of the service will depend on the density of ad hoc network deployment in the area in which the user is located. If there are no nodes to form an ad hoc route from the user to the LI subject, then clearly the system will not work.

8.5 Requirements and Architecture

This section outlines the requirements on devices that are to be involved in the provision of this service, and introduced two possible scenarios — an infrastructure based scenario, and a pure ad hoc network based scenario.

8.5.1 Requirements

Every device which is to be located using the scheme described here must be capable of broadcasting data. Any device that the user wants to use as an LI requestor will need to store data about the devices that the user may wish to locate. All devices should be able to operate within the ad hoc network using the existing routing protocols.

The LI requestor will need to have a measure of location-awareness, i.e. to have some information about its current location. This is necessary in order for the LI requestor to be able to provide a user-accessible interpretation of the LI it receives regarding the LI subject. The location-awareness may be absolute and precise, e.g. as provided by a GPS receiver, or it may only be relative to some other device.

The LI requestor will also need a user interface capable of conveying LI to the user. This might be achieved using a compass style direction indicator, or a more sophisticated graphical display. Current mobile phones and PDAs will clearly be adequate in this respect.

The requirements on the devices to be located will depend on the environment of use, and we now describe some possible usage scenarios.

8.5.2 Infrastructure Based Tracking

The first scenario makes use of an existing location infrastructure, and we use the setting of a car park. We suppose that the LI subject is the user's car. The car park is divided into zones, and each zone has a beacon device. These beacons simply transmit their identities either periodically, or upon request (which may be authenticated). Each car has a means of receiving and processing information from the beacons, and is also capable of acting as a member of an ad hoc network.

When a user needs to find his car, he uses his mobile phone to form an ad hoc network with all the ad hoc enabled devices, in this case including at least some of the cars in the car park. The mobile phone broadcasts a $TrackingRequest$ which contains the mobile phone identifier, the identifier of the car and, optionally, the zone in which the user is located. This $TrackingRequest$ is propagated throughout the ad hoc network until the request reaches the car. The car unicasts a $DirectionReply$ back to the mobile phone, containing the identifier of the beacon closest to the target car (which might be determined on the basis of signal strength). When the mobile phone receives this, it could show the user a map of the car park and where the car is located. This map could be downloaded onto the mobile phone when the user enters the car park as part of a location-based service. If this is not possible, then the zone identifier could be displayed and a map could be provided at frequent points on the walls of the car park.

8.5.3 Ad Hoc Tracking

In the second scenario, again concerned with locating a car, we suppose that either one or both of the mobile phone and car is not within range of a location infrastructure device. Here we need an alternative means of relaying the LI to the user. As the LI subject cannot be sure whether the LI requestor is linked to an infrastructure location node, it has to provide LI which is not dependent on the infrastructure. If the LI subject has a GPS device installed, it could send its location coordinates as LI. However, if the LI requestor has no map then this may be useless information. Even with a map, the user may still be confused as to the direction in which to move. Information which would be more useful to the user is a direction and, possibly, a distance. This could be relayed to the user in the form of a graphical compass arrow and an estimated distance.

The LI could thus be relayed in one of the following ways:

1. *Physical route method:* This is similar to how a source route in the DSR protocol is constructed. Here, each intermediate node appends the direction from which it received the *DirectionReply* message to the Physical route field of the packet. This provides the LI gatherer with a sequence of directions to follow in order to reach the LI subject.

2. *Periodic beaconing method:* Every intermediate node receiving a *DirectionReply* message periodically broadcasts the identifier of the LI subject, and the direction from which the *DirectionReply* was received. Thus, as the LI gatherer moves within transmission range of an intermediate node, it can pick up the beacon. This is particularly useful when the LI subject is mobile and periodically sends a *DirectionReply* message to indicate its new location. Also, the hop count may be included in the *DirectionReply* message, indicating how many hops away the LI subject is. Thus, as the hop count in the received *DirectionReply* messages decreases, the LI gatherer can determine that it is getting closer to the LI subject.

8.5.4 Other Scenarios

As outlined in section 8.1, the service is applicable to many other scenarios. A warehouse stock management scenario is ideally suited for the infrastructure tracking service, as mobility will be low and a limited infrastructure is very feasible. Here, RFID tags [134] may be used for each item, and their presence could be picked up and collated by the infrastructure nodes.

Finally, military scenarios are likely to benefit from an ad hoc tracking service in environments where infrastructure is likely to be limited or even non-existent.

8.6 Location Technology Overview

We now describe a variety of location determining techniques which could be used to help deliver the desired service. With each scheme we provide a discussion of its relative advantages and disadvantages in the context of the location service. The likelihood is that, in order to provide an accurate service, more than one technology will need to be combined. For example, if parts of an ad hoc network contain GPS capable devices, then the LI provided from these devices could be used with other LI to provide a more accurate location service.

8.6.1 The GPS Method

If both the LI requestor and LI subject can discover their coordinates using GPS, then the LI subject can send its coordinates to the LI gatherer via a *DirectionReply* message. The LI requestor can readily combine the received coordinates with its own coordinates to calculate the distance and direction of the LI subject.

However, if the LI gatherer does not know the direction in which it is pointing, it will not be able to convey this direction information in a useful form to the user. Determining the orientation of the LI requestor will require the device to move. In such a case the device could use its new coordinates and the previous coordinates to display a direction for the user to move towards the LI subject. This feature exists

with many current GPS devices [146]. However, the disadvantage of using GPS is that it is very inaccurate indoors. Hence, using GPS would not be suitable for the warehouse scenario. Also, in this situation, GPS may not be accurate enough to pinpoint individual items.

However, the car parking scenario could readily use the GPS method, as many cars are equipped with GPS capable devices. The locating device does not need to be GPS capable, as the cars themselves can calculate a relative location for the LI gatherer to use. Of course, because GPS is inaccurate indoors, this scenario would only apply to outdoor car parks.

Military scenarios could use the Precise Positioning Service (PPS) [146], which give an even greater accuracy than the civilian enabled Standard Positioning Service (SPS) [146].

However, there are many disadvantages to using GPS, as has been widely discussed [118]. The relatively high cost of equipment and the lack of accuracy indoors are among the main issues with using GPS [67].

8.6.2 The Smart Antenna Method

If a mobile device is equipped with a directional antenna, then this could be used to help provide the location discovery service. Ramanathan [116] gives an overview of the possible uses of such antennas in ad hoc networks, along with a discussion of possible advantages and disadvantages. Directional antennas can be used to help provide the service described in this chapter through direction of arrival (DOA) techniques. DOA techniques attempt to determine the direction from which a radio transmission has been received.

If the wireless device can determine from which direction a transmission was received, then this information can be included in the *DirectionReply* messages. If a device is also fitted with an electronic compass, then the *DirectionReply* could also include a compass heading.

The use of smart (directional) antennas would allow the service to be provided in the absence of any pre-existing location measurement infrastructure. Line of sight problems can be overcome, since the path from the LI gatherer to the LI subject can go around obstacles.

The main disadvantage of using directional antennas for wireless communication is the size and relative cost. However, as Ramanathan [116] states, antenna size is decreasing as technology becomes more advanced.

8.6.3 DOA for Omnidirectional Antennas

A possible DOA technique for devices with onmidirectional antennas is as follows. This idea uses the same techniques that the human brain uses to determine the direction from which sound originates. A device would need two aerials spaced as widely as possible. As the device receives a reply it can determine the DOA of the *DirectionalReply* by measuring the differences between the strengths, frequencies and/or times of the two received signals.

Harter et al. [67] apply a similar technique by measuring the time difference between two 'bats' in order to determine the orientation of an object. They state that the greater the distance between the 'bats' the better the orientation measure.

The Cricket compass scheme [114] uses the differences in distance between sensors on a device to determine orientation. However, the authors state that, with current technology, this cannot be achieved reliably, and so they outline other techniques to improve the accuracy of their system.

8.7 Security Requirements

This section is dedicated to exploring the security requirements of an ad hoc network based location tracking service. The security concerns lie largely with privacy and authentication. In a hostile environment, where there may exist many nodes from multiple domains, it is possible that some nodes are not trustworthy.

An unauthorised node is defined as one which is not authorised to view or infer LI regarding a LI subject. One possible security requirement is that it should not be possible for an unauthorised node to link LI subjects and LI gatherers. Doing so compromises the privacy of both the target device and the LI gatherer. For example, if an unauthorised node discovers that LI gatherer A is requesting the location of LI subject B, then it may be able to deduce that A is related to B. In the car park scenario, the unauthorised node could deduce that a particular car is owned by a certain person.

Another security requirement may be that it should not be possible for an unauthorised node to acquire information linking LI subject and LI requestors by posing as an LI subject, posing as a locating device, passively eavesdropping on communications, or by subverting a valid LI subject or LI requestor. Due to the likely mobile nature of devices used in ad hoc networks, the probability that it may be lost or stolen is greater than with desktop computers, for example. For this reason, particular attention must be paid to preventing access to information in compromised devices.

Security requirements may also extend to preventing unauthorised nodes learning of a device's presence. Not only should it be impossible for an unauthorised node to find the precise location of nodes, it should also not be possible for them to learn of the existence of such nodes.[8.4]

If a traditional authentication mechanism is being used, then node existence may be inferred by receiving messages which deny access to LI. Using the military scenario as an example, when an enemy receives a message stating that access to some LI is denied, then they may still deduce a node exists in the direction from which the signal was received, which is an undesirable property. In this case, anonymity may also be a requirement.

Authentication mechanisms should be in place to prevent unauthorised nodes from discovering LI by accessing the location discovery service. The prevention of denial of service attacks is also a potential requirement which may be of particular importance in the military scenario; in this latter case, if LI is denied to an LI requestor, then

[8.4] This contrasts with some sensor networks, where it is the task of the sensor network to detect the existence of foreign nodes.

the LI subject may be incorrectly assumed to be an enemy device.

An unauthorised node should not be able to acquire LI by replaying intercepted messages. This means that replay prevention is required.

Finally, it is also prudent to mention user acceptance of location systems, since this is both an important issue in its own right and a driver for security in such schemes. Some techniques, such as the location tracking and prediction service proposed by Liu, Bahl and Chlamtac [91], could arouse opposition, and a low uptake of the service could potentially result. Such a reaction could occur despite the fact that, as in this latter case, the service could enhance connection reliability by managing cell handoffs more effectively.

8.8 Security Solutions

Securing the routing of protocol messages should be the responsibility of the underlying ad hoc routing protocol. For example, the routing protocol should provide availability, so that if a route exists between an LI requestor and an LI subject, then the service should be successful in sending the location discovery service messages between the two. There is significant existing work on this topic (see, for example, [151]), so we do not address this issue further here.

In Chapter 5, we discuss the LIPA model. There are a number of ways that the requirements for this model are different in an ad hoc network. Observe that this model is designed for use in an infrastructure-based network. In particular, the LI gatherer in the model was a trusted party, and was responsible for creating an LI token. In an ad hoc network, the LI gatherer may be a combination of nodes, some of which may be not be trusted. We also stated that the LIPA entity can be both physically and logically separated from any other entity. Of course, in an ad hoc network, where there is no infrastructure, this poses a problem. The entity requesting LI, i.e. the LI requestor, simply wishes to locate the LI subject. This eliminates the requirements for constraints in this model. The LIPA model is primarily designed to limit LI distribution to service providers.

When deployed in an ad hoc network, we therefore suppose that the LIPA entity simply acts as a device to authenticate LI requestors. If the LI requestor is authorised to have an LI subject's LI, then this information can be encrypted and sent back to the LI requestor. We envisage two different scenarios for gathering LI. One is when the LI gatherer is a single entity, and the other is when the LI gatherer is a number of different nodes in an ad hoc network. These nodes may lie between the LI requestor and the LI subject.

Onion routing techniques can be used to prevent eavesdroppers from finding information about the source and destination of data that is transmitted. When the LI gatherer is a combination of nodes, each of these nodes could direct the LI requestor to the next node. Each node can provide the direction to the next node. These nodes will have been granted permission to pass on information about the direction to the next node. If onion routing techniques are used to hide the LI data, each node will only know the direction to the next node.

The control of access to LI obtained by the LI gatherer is clearly an important issue. Conventionally, this requires the use of an authentication mechanism. This might be possible in the car park scenario, so that only LI subjects within the car park, and possibly only those LI subjects which have subscribed to the service, can use the service provided by the LI gatherer. However, where an infrastructure does not exist it will be difficult to provide an authentication mechanism; indeed, one of the advantages of the proposed service is that it can be provided by a set of ad hoc devices which meet for the first time. One possible solution would involve device manufacturers collaborating to support a key management infrastructure for all mobile devices. Such a scheme may require standardisation for such devices.

If the location messages contain map coordinates or zone identifiers, as in the infrastructure based tracking service (see section 8.5), then these messages should be encrypted to provide confidentiality. Integrity checks could also be provided. Such pairing of devices is not uncommon. An example of a scenario where this occurs is when devices in a Bluetooth network create a security association through pairing.

Our service has the advantage over conventional ad hoc network security schemes that we can assume that the LI requestor and the LI subject have a security associ-

ation. This could, for example, have been set up by the owner of the devices. Thus both symmetric and asymmetric cryptography could be used to provide end-to-end protection.

An important threat arises from compromised devices. There does not appear to be a feasible means for LI subjects to determine whether or not an LI requestor has been compromised. Users may have the option of 'locking out' LI requestors which are believed to be compromised. Unfortunately determining the LI requestors which are thought to be compromised still remains a problem. Thus, the security of LI gatherers will depend on the security properties of the devices themselves, e.g. whether they incorporate password-authentication mechanisms and the level of physical security provided.

The possible physical compromise of LI subjects to reveal their secret keys poses a more interesting problem. Again, physical security measures would ideally be provided; however, LI subjects may sometimes be too small (and low cost) to allow high level physical security protection. One solution is to use short-lived keys to limit the window of opportunity afforded to an attacker who retrieves compromised keys.

8.9 Related Work

Much research has already been performed in this area, and many schemes have been proposed that offer a similar location tracking service using different technologies. However, not much has been written about the associated security issues. We now give an overview of the advantages and disadvantages of the various existing techniques, and also, where relevant, highlight the security concerns which have been raised. Hightower and Borriello [72] provide a taxonomy of location systems and give a survey of current research.

The 'Active Badge' location system [149] provides a similar service, but in an indoor environment. This system relies on infra-red technology, where sensors detect periodic signals emitted by 'Active badges'. These signals are collated and processed by a central server. This information is either relayed via a desktop application, or

used to automate the routing of telephone calls in a Public Branch Exchange (PBX) telephony network. The 'Active badge' successor, the 'BAT' system [67], uses ultrasound techniques. The main difference between our scheme and the 'Active Badge' system is that the latter depends on a infrastructure backbone and a central server, the presence of which cannot be assumed in an ad hoc network. Also, Hightower and Borriello [72] highlight the limitations of using infrared and ultrasound technology. Want et al. [149] discuss the privacy issues arising from use of the Active Badge system. In particular, they consider concerns about the misuse of LI and giving users the right not to wear Active Badges.

The 'EasyLiving' tracking system [86] uses computer vision techniques to track the location of people in an indoor intelligent environment. Stereo camera images in the room are analysed for 'blobs', which are used to form the shape of a human figure, and additional information such as colour histograms are used for identification purposes. So, while the application is similar to ours, 'EasyLiving' provides a more specialised system, which again relies on an indoor infrastructure and probably also does not scale well. Hightower and Borriello [72] again provide more technical insight into the difficulties of using vision location systems.

Typically, location-based schemes are designed for use in an environment where devices can locate their own position, either on a map or by learning geographical coordinates. Wireless LANs (WLAN) have become extremely popular over the past few years, with IEEE 802.11 standards emerging as the dominant technology. Tao et al. [137] is one of many examples of research aimed at achieving location discovery of WLAN devices using radio frequency techniques (see also [9, 131]).

The Tao et al. system is designed for use indoors, and requires the presence of a number of fixed wireless access points. Thus an infrastructure is required, with a central server that controls location sensing. Initially, an offline training phase takes place, in which the server builds a conditional probability distribution in order to determine the likely future location of an LI subject. This is achieved by the use of 'snoopers', which may be fixed access points or mobile laptops. The server can then use the conditional probability distribution with Bayesian inference to determine the location of a WLAN device. The authors explain how to use their system to locate rogue machines which are attempting to gain unauthorised access to the

building network. Whereas previous systems are vulnerable to a rogue attacker varying their broadcast power to remain undetected, Tao et al. [137] claim that their algorithms are not vulnerable to such an attack. However, they do highlight several security concerns for their system, including the possibility that an attacker could set up its own 'snoopers' and locate WLAN devices. This problem arises because wireless broadcast devices cannot choose who will receive their signals. War-driving is given as an example threat. As previously, the system has implementation problems compared to the solution we propose. There is a reliance on a centralised and fixed infrastructure. In addition, most WLAN schemes need a training period and so would not be usable in a dynamic environment. Finally, Hightower and Borriello [72] point out that WLAN technology may not be available on smaller devices.

Smailagic and Kogan [131] present a 'Portable Help Desk', so that users in a university campus can locate other users and see their contact information. They also discuss how user location privacy can be provided through the use of time scheduled rules. Each rule determines the visibility[8.5] of the user during a certain period of time. From a survey of users which used the 'Portable Help Desk', they reveal that users are unwilling to engage in setting up offline communications channels, preferring interactive access control or just letting anyone communicate. With respect to who can see a user online in the system, again users were willing for anyone to see them, but here more users preferred offline access control.

The emergence of sensor networks has also provided a catalyst for location detection. Doherty, Pister and El Ghaoui [39] present a position estimation method for ad hoc sensor networks, in which all sensors send their connection information to a central computer that calculates the positions of every sensor in the network. In the context of our location discovery service, the central computer would be the LI gatherer. However, this system requires that the intermediate nodes must reveal their location, which may be unacceptable for privacy reasons.

Capkun, Hamdi and Hubaux [22] introduce the Self-Positioning Algorithm (SPA), where a node can determine its own relative position in the ad hoc network. This might be useful to allow the LI gatherer and LI subject to individually discover

[8.5] Visible to all, Invisible to some, Visible to some or Invisible to all.

their own relative locations. The LI subject can then sends its coordinates to the LI gatherer. However this system suffers from the fact that cooperation is needed between a number of nodes, where the accuracy of the system increases as the number of involved nodes rises; hence privacy issues again arise. Also, all nodes have to align their 'Network Coordinate System' to a 'Local Reference Group', and the extra overhead of control messages may be too great for a wireless network. This system would also be very vulnerable to attack. Finally, there are issues with relying on Time of Arrival techniques to determine distance.

Priyantha, Chakraborty and Blakrishnan [113] introduce a decentralised scheme based on both radio frequency and ultrasound techniques, where there is no central database of control. The 'Cricket' scheme is designed to give location support to services made available to a user. They claim that their solution is scalable and enhances user privacy. The system uses a beacon infrastructure, and thus this technology may be utilised in the car park scenario. The authors also introduce the concept of a map server, by which nodes in the network can discover what services are available by downloading an active map.

In [114], Priyantha et al. extend the 'Cricket' scheme to create an electronic compass called 'Wayfinder', and a service discovery scheme called 'viewfinder'. The 'Wayfinder' is identical in purpose to our application. However, because of its reliance on the beacon infrastructure, and the fact that a map has to be pre-installed on the devices, our ad hoc network based solution is more dynamic and easily adaptable.

Much work has been performed on location services for cellular technology such as CDMA and GSM, and Zagami et al. [152] provide a useful overview. This work has been motivated by the E911 FCC ruling that all specialised mobile radio and personal communication systems making 911 emergency calls should be automatically located to within 125 metres. Numerous commercial GSM location services are already available in Europe. These location services operate by modifying the base stations to use time of arrival (TOA) and angle of arrival (AOA) techniques. The main advantage is that no modification of existing mobile handsets is required. The main disadvantage is that data from three or more base stations are needed to determine a position, which means that the system is sometimes not available. Zagami

et al. [152] also suggest other applications for a location discovery service, such as tracking missing/lost Alzheimer's patients and also tracking tagged criminals. Many of these schemes could be used in conjunction with our location discovery tracking service to provide a comprehensive set of location functions to the user. For example, a user with a GSM phone in a car park could call the phone number of a stationary locating device in the car park. This LI gatherer could then perform the car location service on behalf of the user. Cellular location techniques could be used to track the user, and the ad hoc location technique could be used to track the car. These could then be combined to direct the user. However, a coordinated effort is required in order to make the different systems interoperate in this way.

Finally, Bauer, Becker and Rothermel [13] present a location modelling language, which could be useful for implementing applications which use our scheme.

8.10 Future Work

The next logical step in this research would involve an investigation into the messages transferred in this scheme. This would enable a quantitative measurement to be made of the cost of deployment of such a scheme. Further, simulation of this scheme would enable an analysis of its efficiency.

Research into the underlying technologies which would enable this service to function, such as those touched upon in section 8.6, could create greater efficiency in location calculation. This, of course, would also improve the general efficiency of this scheme.

The security requirements, detailed in section 8.7 above, reveal several important issues regarding the privacy of LI, and some of these have been discussed in section 8.8. However, much more research is needed on issues such as key management for the use of public key cryptography in a mobile ad hoc setting.

8.11 Summary and Conclusions

We have described how a locating service may be useful in a variety of scenarios, and we have introduced the notion of providing such a service using ad hoc network routing principles. In particular, we have shown the requirements on an infrastructure for such a service, and evaluated ways in which this may be implemented using a variety of different technologies. Security requirements for the deployment of such a service, with a focus on authentication and privacy, have been discussed. An overview of solutions proposed by other authors has also been provided. Finally, we examined future directions for this research.

Chapter 9

Conclusions

Contents

9.1 Summary and Conclusions 152
9.2 Suggestions for Further Work 153

This chapter summarises this thesis, provides conclusions and gives an overview of the achievements of this thesis. It also discusses further work and possible future directions for this research.

9.1 Summary and Conclusions

In this section a summary of the work in this thesis is provided.

We began with an analysis of user requirements for privacy when using a location-based service, and then described a means of meeting these requirements in the form of constraints in Chapter 4. However, we also observed that the constraints themselves pose a threat to user privacy. Additionally, simply describing user requirements for privacy does not ensure that information consumers will abide by them.

In Chapter 5, we proposed a means of overcoming the limitations of providing constraints by introducing a LIPA framework. The LIPA is a trusted third party which analyses an LI token containing constraints, and then provides LI to permitted entities. Before LI is provided to the entity, the constraints are removed from the LI token. This ensures that additional private information is not disseminated.

Chapter 6 discussed TC and the TCG specifications, and considered how TC functionality could be used to gain assurance regarding the integrity of software on a remote platform. The purpose of this investigation was to tackle the the problem of controlling LI once it had been disseminated, which was highlighted in Chapter 5. We discussed how the integrity of software on the platform to which personal information is being sent could be verified using TC, providing assurance of its trustworthiness.

In Chapter 7, we described how LI constraints can be modelled using the UCON_{ABC} model. This provided a formal structure to describe the control of LI.

Finally, Chapter 8 described a location discovery service using ad hoc network routing principles. We investigated the requirements on an infrastructure for such a service, and how such a service could be implemented using a variety of technologies. We also discussed the security requirements for the deployment of such a service.

9.2 Suggestions for Further Work

Numerous suggestions for further work have been made throughout this thesis. This section summarises some of the most noteworthy examples.

Chapter 4 discussed the use of constraints in the context of controlling LI. However, the use of constraints does not need to be limited only to LI. Constraints could also be used in a variety of other contexts, for example temperature or physical events. Additionally, privacy may be considered with a variety of entities in mind: for example, the focus of the scheme could be a corporation rather than an end user. Constraints used for a variety of entities or contexts may also have different security requirements. This may provide new directions for research.

A formal language to describe LI constraints could be used in a variety of different solutions to the LI privacy problem. This language should be extensible, so it can provide constraints not only for LI scenarios, but also for different entities and in multiple contexts.

It is also difficult to ensure that entities abide by the constraints which are set by the LI subject, and to prove when the constraints have been abused. Finding ways to address such issues is an important research challenge.

Concerns outlined in Chapter 6 include the difficulty of establishing the potential uses of information by a service provider. A possible solution to this problem is for trusted sources to host lists with information about service providers. This information could be based on the reputation of the service provider. However, establishing such lists is a difficult problem. There is no way of establishing how information will be handled once a service is provided. Enabling such features may provide interesting research possibilities.

Chapter 6 also discusses a mechanism to establish the trustworthiness of software used to handle personal information. This is also, in itself, a difficult task. Third party testing laboratories should ideally conform to some agreed procedures which can provide the end user with a certain level of trust. Such standardised testing may provide end users with a minimum level of trust in software used to manage their information. Establishing the basis for such standardised testing for software may also provide directions for fruitful research possibilities.

In Chapter 7 we examined the problem of providing privacy for LI in the UCON_{ABC} model. Various limitations of the model meant that we could not express controls which apply when access to LI is denied. Further research on this topic could investigate means to describe such scenarios.

Chapter 8 investigated a location discovery service within an ad hoc network. Further research in this area could investigate the volume and size of messages transferred in an implementation of the scheme. Additional analysis of this scheme could also investigate its efficiency. Solutions to the security requirements described in the Chapter may also reveal important issues regarding the privacy of location information. For example, issues such as key management for the use of public key cryptography in a mobile ad hoc setting may itself be an interesting direction for further research.

Bibliography

[1] 104th Congress, Senate and House of Representatives of the United States of America. *Telecommunications Act of 1996*, 1996.

[2] 3rd Generation Partnership Project. *3GPP TS 03.71 V8.7.0 Technical Specification Group Services and System Aspects; Location Services (LCS); (Functional description) Stage 2 (Release 1999)*, September 2002.

[3] 3rd Generation Partnership Project. *Characteristics of the USIM application*, v7.7.0 edition, November 2006.

[4] L. Ackerman, J. Kempf, and T. Miki. Wireless location privacy: Law and policy in the US, EU and Japan. ISOC Member Briefing 15, Internet Society, November 2003.

[5] A. Adams. Users' perception of privacy in multimedia communication. In *CHI '99 extended abstracts on Human factors in computing systems*, pages 53–54. ACM Press, New York, NY, USA, May 1999.

[6] H. Alvestrand. A mission statement for the IETF. RFC 3935, Internet Engineering Task Force, October 2004.

[7] W. A. Arbaugh, D. J. Farber, and J. M. Smith. A secure and reliable bootstrap architecture. In *Proceedings of the 1997 IEEE Symposium on Security and Privacy*, pages 65–71, 1997.

[8] J. Arkko, V. Devarapalli, and F. Dupont. Using IPsec to protect mobile IPv6 signaling between mobile nodes and home agents. RFC 3776, Internet Engineering Task Force, June 2004.

[9] P. Bahl and V. N. Padmanabhan. RADAR: An in-building RF-based user location and tracking system. In *Proceedings of the Nineteenth Annual Joint*

Conference of the IEEE Computer and Communications Societies (INFOCOM 2000), volume 2, pages 775–784, March 2000.

[10] B. Balacheff, L. Chen, S. Pearson, D. Plaquin, and G. Proudler. *Trusted computing platforms: TCPA technology in context*. Hewlett-Packard professional books. Prentice-Hall, Englewood Cliffs, NJ, USA, 2002.

[11] B. Balacheff, L. Chen, D. Plaquin, and G. Proudler. A trusted process to digitally sign a document. In *Proceedings of the 2001 Workshop on New Security Paradigms (NSPW '01)*, pages 79–86, September 2001.

[12] L. Barkhuus and A. Dey. Location-based services for mobile telephony: a study of users privacy concerns. In G. W. M Rauterberg, M. Menozzi, and J. Wesson, editors, *Human-Computer Interaction INTERACT '03: IFIP TC13 International Conference on Human-Computer Interaction*. IOS Press, Amsterdam, Netherlands, September 2003.

[13] M. Bauer, C. Becker, and K. Rothermel. Location models from the perspective of context-aware applications and mobile ad hoc networks. *Personal and Ubiquitous Computing*, 6:322–328, 2002.

[14] R. Beckwith. Designing for ubiquity: The perception of privacy. *IEEE Pervasive Computing*, 2(2):40–46, April 2003.

[15] D. E. Bell and L. J. La Padula. Secure computer systems: Unified exposition and multics interpretation. Technical Report ESD-TR-75-306, The Mitre Corporation, March 1976.

[16] B. Bellur and R. Ogier. A reliable, efficient topology broadcast protocol for dynamic networks. In *Proceedings of the Eighteenth Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM '99)*, volume 1, pages 178–186. IEEE Press, Piscataway, NJ, USA, March 1999.

[17] A. R. Beresford and F. Stajano. Location privacy in pervasive computing. *IEEE Pervasive Computing*, 2(1):46–55, 2003.

[18] A. R. Beresford and F. Stajano. Mix zones: User privacy in location-aware services. In *Proceedings of the Second IEEE Annual Conference on Pervasive Computing and Communications Workshops*, pages 127–131. IEEE Computer Society Press, Los Alamitos, CA, USA, March 2004.

[19] T. Bray, J. Paoli, C. M. Sperberg-McQueen, E. Maler, and F. Yergeau. Extensible markup language (XML) 1.0 (third edition). W3C recommendation, World Wide Web Consortium, February 2004.

[20] A. Burak and T. Sharon. Analyzing usage of location based services. In *Extended abstracts on Human factors in computing systems (CHI '03)*, pages 970–971. ACM Press, New York, NY, USA, April 2003.

[21] S. Byers and D. Kormann. 802.11b access point mapping. *Communications of the ACM*, 46(5):41–46, May 2003.

[22] S. Capkun, M. Hamdi, and J. Hubaux. GPS-free positioning in mobile ad-hoc networks. *Cluster Computing Journal*, 5(2):157–167, 2002.

[23] D. Chaum. Achieving electronic privacy. *Scientific American*, 267:96–101, 1992.

[24] D. L. Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24(3):84 – 90, February 1981.

[25] L. Chen and S. Pearson. A trusted biometric system. Technical Report HPL-2002-185, HP Laboratories Bristol, July 2002.

[26] B. Chor, E. Kushilevitz, O. Goldreich, and M. Sudan. Private information retrieval. *Journal of the ACM*, 45(6):965–981, November 1998.

[27] Cingular Wireless II LLC, http://www.cingular.com/mmode/mmode_net. *Features and Services Information for Former AT&T Wireless Users*, September 2005.

[28] Cingular Wireless II LLC, http://www.cingular.com/privacy/privacy_policy. *Privacy Policy*, September 2005.

[29] P. C. Clark and Lance J. Hoffman. BITS: A smartcard protected operating system. *Communications of the ACM*, 37(11):66–70, November 1994.

[30] T. Clausen, G. Hansen, L. Christensen, and G. Behrmann. The optimized link state routing protocol, evaluation through experiments and simulation. In *Proceedings 4th International Symposium on Wireless Personal Multimedia Communications*, pages 841–846. IEEE Press, Piscataway, NJ, USA, September 2001.

[31] L. Cranor, B. Dobbs, S. Egelman, G. Hogben, J. Humphrey, M. Langheinrich, M. Marchiori, M. Presler-Marshall, J. Reagle, M. Schunter, D. A. Stampley, and R. Wenning. The platform for privacy preferences. W3C recommendation, World Wide Web Consortium, November 2006.

[32] L. F. Cranor and B. A. La Macchia. Spam! *Communications of the ACM*, 41(8):74–83, August 1998.

[33] J. Cuellar, J. Morris, D. Mulligan, J. Peterson, and J. Polk. Geopriv requirements. RFC 3693, Internet Engineering Task Force, February 2004.

[34] M. Danley, D. Mulligan, J. Morris, and J. Peterson. Threat analysis of the geopriv protocol. RFC 3694, Internet Engineering Task Force, February 2004.

[35] S. Deering and R. Hinden. Internet protocol, version 6 (ipv6) specification. RFC 2460, IETF, December 1998.

[36] D. E. Denning and P. F. MacDoran. Location-based authentication: Grounding cyberspace for better security. In D. E. Denning and P. J. Denning, editors, *Internet Besieged, Countering Cyberspace Scofflaws*, chapter 12, pages 167–174. ACM Press, New York, NY, USA, 2nd edition, February 2001.

[37] A. W. Dent and C. J. Mitchell. *User's Guide to Cryptography and Standards*. Artech House, London, UK, 2004.

[38] R. Dingledine and N. Mathewson. Tor: The second-generation onion router. In *Proceedings of the Thirteenth USENIX security symposium*, pages 303–320. USENIX, Berkeley, CA, USA, August 2004.

[39] L. Doherty, K. S. J. Pister, and L. El Ghaoui. Convex position estimation in wireless sensor networks. In *Proceedings of the Twentieth Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM 2001)*, volume 3, pages 165–1663, April 2001.

[40] R. Droms. Dynamic host configuration protocol. RFC 2131, Internet Engineering Task Force, March 1997.

[41] J. Dunlop, R. C. Atkinson, J. Irvine, and D. Pearce. A personal distributed environment for future mobile systems. In *Proceedings of the IST Mobile and Wireless Communications Summit*, pages 705–709. Instituto de Telecomunicações, Portugal, June 2003.

[42] S. Duri, M. Gruteser, X. Liu, P. Moskowitz, R. Perez, M. Singh, and J.-M. Tang. Framework for security and privacy in automotive telematics. In *Proceedings of the 2nd International Workshop on Mobile Commerce (WMC'02)*, pages 25–32. ACM Press, New York, NY, USA, September 2002.

[43] S. Dusse, P. Hoffman, B. Ramsdell, L. Lundblade, and L. Repka. S/mime version 2 message specification. RFC 2311, IETF, March 1998.

[44] Senator J. Edwards. Location privacy protection act of 2001. Bill S.1167, US Senate, July 11, 2001.

[45] P. Enge and P. Misra. Special issue on global positioning system. *Proceedings of the IEEE*, 87(1):3–15, 1999.

[46] P. England, B. Lampson, J. Manferdelli, M. Peinado, and B. Willman. A trusted open platform. *Computer*, 36(7):55–62, July 2003.

[47] P. England and M. Peinado. Authenticated operation of open computing devices. In L. Batten and J. Seberry, editors, *Proceedings of the 7th Australasian Conference on Information Security and Privacy (ACISP 2002), Melbourne, Australia, July 3-5, 2002*, volume 2384 of *Lecture Notes in Computer Science*, pages 346–361. Springer-Verlag, Berlin, Germany, July 2002.

[48] M. Epstein and S. Vergani. History unwired: mobile narrative in historic cities. In *Proceedings of the working conference on Advanced visual interfaces (AVI '06)*, pages 302–305. ACM Press, New York, NY, USA, May 2006.

[49] A. Escudero-Pascual and G. Q. Maguire Jr. Role(s) of a proxy in location based services. In *Proceedings of the 13th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications*, volume 3, pages 1252–1256. IEEE, September 2003.

[50] Federal Communications Commission. *ORDER DA 02-2423, Revision of the Commission's Rules to Ensure Compatibility with Enhanced 911 Emergency Calling Systems*, 2002.

[51] E. W. Felten. Understanding trusted computing: Will its benefits outweigh its drawbacks? *IEEE Security and Privacy*, 1(3):60–62, May 2003.

[52] W. Ford. *Computer Communications Security*. Prentice-Hall, Englewood Cliffs, NJ, USA, 1994.

[53] D. Fox, J. Hightower, L. Liao, and D. Schulz. Bayesian filtering for location estimation. *IEEE Pervasive Computing*, 2(3):24–33, 2003.

[54] N. Freed and N. Borenstein. Multipurpose internet mail extensions (mime) part one: Format of internet message bodies. RFC 2045, IETF, November 1996.

[55] A. S. Gajparia. On location-based services and the usage control model (extended abstract). In *Western European Workshop on Research in Cryptology*, pages 74–77. WEWoRC Conference Records, Leuven, Belgium, July 2005.

[56] A. S. Gajparia and C. J. Mitchell. Enhancing user privacy using trusted computing. In C. J. Mitchell, editor, *Trusted Computing*, chapter 8, pages 239–249. IEE, Hertfordshire, UK, 2005.

[57] A. S. Gajparia, C. J. Mitchell, and C. Y. Yeun. Using constraints to protect personal location information. In *Proceedings of the 58th IEEE Vehicular Technology Conference (VTC 2003-Fall)*, volume 3, pages 2112–2116. IEEE Press, Piscataway, NJ, USA, October 2003.

[58] A. S. Gajparia, C. J. Mitchell, and C. Y. Yeun. The location information preference authority: Supporting user privacy in location based services. In S. Liimatainen and T. Virtanen, editors, *Proceedings of Nordsec 2004, the 9th Nordic Workshop on Secure IT systems*, pages 91–96. Helsinki University of Technology, Finland, November 2004.

[59] A. S. Gajparia, C. J. Mitchell, and C. Y. Yeun. Supporting user privacy in location based services. *IEICE Transactions*, E88-B(7):2848–2855, July 2005.

[60] E. Gallery. An overview of trusted computing technology. In C. J. Mitchell, editor, *Trusted Computing*, chapter 3, pages 29–112. IEE, Hertfordshire, UK, 2005.

[61] P. B. Gibbons, B. Karp, Y. Ke, S. Nath, and S. Seshan. Irisnet: An architecture for a worldwide sensor web. *IEEE Pervasive Computing*, 2(4):22–33, October – November 2003.

[62] D. Gollmann. *Computer Security.* John Wiley and Sons, Chicester, UK, 1999.

[63] M. Gruteser and D. Grunwald. Anonymous usage of location-based services through spatial and temporal cloaking. In *Proceedings of the First International Conference on Mobile Systems, Applications, and Services*, pages 31–42. USENIX, Berkeley, CA, USA, May 2003.

[64] Ceki Gulcü and Gene Tsudik. Mixing e-mail with BABEL. In *Proceedings of the 1996 Symposium on Network and Distributed System Security (SNDSS '96)*, pages 2–16. ACM Press, New York, NY, USA, February 1996.

[65] V. Haldar, D. Chandra, and M. Franz. Semantic remote attestation — virtual machine directed approach to trusted computing. In *Proceedings of the 3rd Virtual Machine Research and Technology Symposium*, pages 29–41. USENIX, Berkeley, CA, USA, May 2004.

[66] A. Harter, A. Hopper, P. Steggles, A. Ward, and P. Webster. The anatomy of a context-aware application. In *Proceedings of the 5th Annual ACM/IEEE International Conference on Mobile Computing and Networking (Mobicom 1999)*, pages 59–68. ACM Press, New York, NY, USA, August 1999.

[67] A. Harter, A. Hopper, P. Steggles, A. Ward, and P. Webster. The anatomy of a context-aware application. *Wireless Networks*, 8(2/3):187–197, 2002.

[68] D. N. Hatfield. A report on technical and operational issues impacting the provision of wireless enhanced 911 services. Technical report, Federal Communications Commission, 2002.

[69] M. Hazas and A. Ward. A high performance privacy-oriented location system. In *Proceedings of the First IEEE International Conference on Pervasive Computing and Communications (PerCom'03)*, pages 216–223. IEEE Computer Society Press, Los Alamitos, CA, USA, March 2003.

[70] U. Hengartner and P. Steenkiste. Implementing access control to people location information. In *Proceedings of the Ninth ACM Symposium on Access Control Models and Technologies (SACMAT '04)*, pages 11–20. ACM Press, New York, NY, USA, June 2004.

[71] R. Hes and J. J. Borking. *Privacy Enhancing Technologies: The path to anonymity.* The Hague, The Hague, Netherlands, revised edition, 1998.

[72] J. Hightower and G. Borriello. Location systems for ubiquitous computing. *Computer*, 34(8):57–66, 2001.

[73] IETF Secretariat. Geographic Location/Privacy (geopriv) Charter. Available at http://www.ietf.org/html.charters/geopriv-charter.html, Internet Engineering Task Force, August 2001.

[74] A. Iliev and S. Smith. Protecting client privacy with trusted computing at the server. *IEEE Security and Privacy*, 3(2):20–28, March 2005.

[75] T. Imielinski and J. C. Navas. GPS-based geographic addressing, routing, and resource discovery. *Communications of the ACM*, 42(4):86–92, April 1999.

[76] International Organisation for Standardization. *ISO 7498-2: Information processing systems – Open Systems Interconnection – Basic Reference Model – Part 2: Security Architecture*, 1989.

[77] International Organization for Standardization. *ISO/TC 211/WG 4/PT 19136: Geographic Information -- Geography Markup Language (GML)*, committee draft edition, February 2004.

[78] International Telecommunication Union. *Standard-frequency and time-signal emissions – annex I*, 1986.

[79] J.Hightower and G. Boriello. A survey and taxonomy of location systems for ubiquitous computing. Technical Report UW-CSE 01-08-03, University of Washington, August 2001.

[80] X. Jiang and J. A. Landay. Modeling privacy control in context-aware systems. *IEEE Pervasive Computing*, 1(3):59–93, July 2002.

[81] D. Johnson, D. Maltz, and J. Broch. DSR — The dynamic source routing protocol for multihop wireless ad hoc networks. In C. Perkins, editor, *Ad Hoc Networking*, chapter 5, pages 139–172. Addison-Wesley, 2001.

[82] A. Juels, R. Rivest, and M. Szydlo. The blocker tag: selective blocking of rfid tags for consumer privacy. In *Proceedings of the 10th ACM Conference on Computer and Communications Security (CCS'03)*, pages 103–111. ACM Press, New York, NY, USA, 2003.

[83] E. Kaasinen. User needs for location-aware mobile services. *Personal and Ubiquitous Computing*, 7(1):70–79, May 2003.

[84] S. Kent and R. Atkinson. Security architecture for the internet protocol. RFC 2401, IETF, November 1998.

[85] R. Koodli. Ip address location privacy and mobile ipv6: Problem statement. MIP6 Working Group Internet Draft draft-ietf-mip6-location-privacy-ps-00.txt, Internet Engineering Task Force, October 2005.

[86] J. Krumm, S. Harris, B. Meyes, B. Brummitt, M. Hale, and S. Shafer. Multi-camera multi-person tracking for easyliving. In *Proceedings of the Third IEEE International Workshop onVisual Surveillance*, pages 3–10. IEEE Press, Piscataway, NJ, USA, July 2000.

[87] B. W. Lampson. Protection. *ACM SIGOPS Operating Systems Review*, 8(1):18–24, January 1974.

[88] U. Leonhardt and J. Magee. Towards a general location service for mobile environments. In *Proceedings of Third International Workshop on Services in Distributed and Networked Environments*, pages 43–51. IEEE Computer Society Press, Los Alamitos, CA, USA, June 1996.

[89] U. Leonhardt and J. Magee. Security considerations for a distributed location service. *Journal of Network Systems Management*, 6(1):51–70, March 1998.

[90] X.-Y. Li, C.-X. Shen, and X.-D. Zuo. An efficient attestation for trustworthiness of computing platform. In *Proceeding of the Second International Conference on Intelligent Information Hiding and Multimedia Signal Processing (IIH-MSP 2006), Pasadena, California, USA, December 18-20, 2006*, pages 625–630. IEEE Computer Society Press, Los Alamitos, CA, USA, December 2006.

[91] T. Liu, P. Bahl, and I. Chlamtac. Mobility modeling, location tracking and trajectory prediction in wireless ATM networks. *IEEE Journal on Selected Areas in Communications*, 16(6):922–936, August 1998.

[92] The Local Data Company Limited, http://www.e-street.com. *Retail Intelligence*, April 2006.

[93] R. Mahy. A document format for filtering and reporting location notications in the presence information document format location object. Geopriv Internet-Draft draft-ietf-geopriv-loc-filters-00.txt (work in progress), Internet Engineering Task Force, March 20, 2006.

[94] M. Maxim and D. Pollino. *Wireless Security.* McGraw-Hill/Osborne, 2002.

[95] A. J. Menezes, P. C. van Oorschot, and S. A. Vanstone. *Handbook of applied cryptography.* CRC Press Series on Discrete Mathematics and its Applications. CRC Press, Boca Raton, FL, USA, 1997.

[96] R. P. Minch. Privacy issues in location-aware mobile devices. In *Proceedings of the 37th Annual Hawaii International Conference on System Sciences (HICSS'04) – Track 5.* IEEE Computer Society Press, Los Alamitos, CA, USA, January 2004.

[97] C. J. Mitchell, editor. *Trusted Computing.* IEE, Hertfordshire, UK, 2005.

[98] G. Myles, A. Friday, and N. Davies. Preserving privacy in environments with location-based applications. *IEEE Pervasive Computing*, 2(1):56–64, 2003.

[99] National Institute of Standards and Technology. Specification for the advanced encryption standard (AES). Federal Information Processing Standards Publication 197 (FIPS PUB 197), NIST, November 2001.

[100] J. Park and R. Sandhu. Originator control in usage control. In *Proceedings of the Third IEEE International Workshop on Policies for Distributed Systems and Networks (POLICY'02)*, pages 60–67. IEEE Computer Society Press, Los Alamitos, CA, USA, June 2002.

[101] J. Park and R. Sandhu. Towards usage control models: beyond traditional access control. In *Proceedings of the Seventh ACM Symposium on Access Control Models and Technologies (SACMAT'02)*, pages 57–64. ACM Press, New York, NY, USA, June 2002.

[102] J. Park and R. Sandhu. The UCON_{ABC} usage control model. *ACM Transactions on Information and System Security*, 7(1):128–174, February 2004.

[103] A. Pashalidis and C. J. Mitchell. Single sign-on using trusted platforms. In C. Boyd and W. Mao, editors, *Procceedings of the 6th International Conference*

on Information Security (ISC 2003) Bristol, UK, October 1-3, 2003, pages 54–68. Springer-Verlag, Berlin, Germany, 2003.

[104] M. Peinado, Y. Chen, P. England, and J. Manferdelli. NGSCB: A trusted open system. In H. Wang, J. Pieprzyk, and V. Varadharajan, editors, *Proccedings of the 9th Australasian Conference on Information Security and Privacy (ACISP 2004), Sydney, Australia, July 13-15, 2004*, volume 3108 of *Lecture Notes in Computer Science*, pages 86–97. Springer-Verlag, Berlin, Germany, July 2004.

[105] M. Peinado, P. England, and Y. Chen. An overview of NGSCB. In C. J. Mitchell, editor, *Trusted Computing*, chapter 4, pages 115–142. IEE, Hertfordshire, UK, 2005.

[106] C. Perkins and E. Royer. *The Ad Hoc On-Demand Distance-Vector Protocol*, chapter 6, pages 173–219. Addison-Wesley, 2001.

[107] C. E. Perkins. *Ad Hoc Networking*. Addison-Wesley, Boston, MA, USA, 2001.

[108] J. Peterson. A presence architecture for the distribution of geopriv location objects. RFC 4079, Internet Engineering Task Force, July 2005.

[109] J. Peterson. A presence-based geopriv location object format. RFC 4119, Internet Engineering Task Force, December 2005.

[110] J. Peterson. A presence-based GEOPRIV location object format. Geopriv Internet-Draft draft-ietf-geopriv-pidf-lo-03.txt (work in progress), Internet Engineering Task Force, 2004 September 9,.

[111] A. Pfitzmann and M. Köhntopp. Anonymity, unobservability, and pseudonymity — a proposal for terminology. In H. Federrath, editor, *Proceedings of Designing Privacy Enhancing Technologies: International Workshop on Design Issues in Anonymity and Unobservability, Berkeley, CA, USA, July 2000*, volume 2009 of *Lecture Notes in Computer Science*, pages 141–160. Springer-Verlag, Berlin, Germany, 2001.

[112] J. Polk, J. Schnizlein, and M. Linsner. Dynamic host configuration protocol option for coordinate-based location configuration information. RFC 3825, Internet Engineering Task Force, 2004 July.

[113] N. Priyantha, A. Chakraborty, and H. Balakrishnan. The cricket location support system. In *Proceedings of the 6th Annual International Conference on*

Mobile Computing and Networking (MobiCom'00), pages 32–43. ACM Press, New York, NY, USA, August 2000.

[114] N. Priyantha, A. Miu, H. Balakrishnan, and S. Teller. The cricket compass for context-aware mobile applications. In *Proceedings of the 7th Annual International Conference on Mobile Computing and Networking (MobiCom'01)*, pages 1–14. ACM Press, New York, NY, USA, July 2001.

[115] N. B. Priyantha, A. Chakraborty, and H. Balakrishnan. The cricket location-support system. In *Proceedings of the 6th Annual International Conference on Mobile Computing and Networking (MobiCom '00)*, pages 32–43. ACM Press, New York, NY, USA, August 2000.

[116] R. Ramanathan. On the performance of ad hoc networks with beamforming antennas. In *Proceedings of the 2nd ACM International Symposium on Mobile Ad Hoc Networking & Computing (MobiHoc '01)*, pages 95–105. ACM Press, New York, NY, USA, October 2001.

[117] C. Randell and H. L. Muller. Low cost indoor positioning system. In G. D. Abowd, B. Brumitt, and S. Shafer, editors, *Proceedings of the Third International Conference on Ubiquitous Computing (Ubicomp 2001), Atlanta, Georgia, USA, September 30 - October 2, 2001*, pages 42–48. Springer-Verlag, Berlin, Germany, September / October 2001.

[118] J. Reed, K. Krizman, B. Woerner, and T. Rappaport. An overview of the challenges and progress in meeting the e-911 requirement for location service. *IEEE Communications Magazine*, 36(4):30–37, April 1998.

[119] C. Rigney, A. Rubens, W. Simpson, and S. Willens. Remote authentication dial in user service (RADIUS). RFC 2138, Internet Engineering Task Force, April 1997.

[120] R. L. Rivest, A. Shamir, and L. M. Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Communications of the ACM*, 26(1):96–99, January 1983.

[121] J. Rosenberg. Presence authorization rules. Internet-Draft draft-ietf-simple-presence-rules-04.txt, Internet Engineering Task Force, October 2005.

[122] A.-R. Sadeghi and C. Stüble. Property-based attestation for computing platforms: caring about properties, not mechanisms. In *Proceedings of the 2004 workshop on new security paradigms (NSPW '04)*, pages 67–77. ACM Press, New York, NY, USA, September 2004.

[123] R. Sandhu and J. Park. Usage control: A vision for next generation access control. In V. Gorodetsky, L. J. Popyack, and V. A. Skormin, editors, *Proceedings of the Second International Workshop on Mathematical Methods, Models, and Architectures for Computer Network Security (MMM-ACNS 2003), St. Petersburg, Russia, September 21-23, 2003*, volume 2776 of *Lecture Notes in Computer Science*, pages 17–31. Springer-Verlag, Berlin, Germany, September 2003.

[124] N. Sastry, U. Shankar, and D. Wagner. Secure verification of location claims. In *Proceedings of the 2003 ACM Workshop on Wireless Security (WiSe '03)*, pages 1–6. ACM Press, New York, NY, USA, July 2001.

[125] H. Schulzrinne. Dynamic host configuration protocol (DHCPv4 and DHCPv6) option for civic addresses configuration information. Geopriv Internet-Draft draft-ietf-geopriv-dhcp-civil-09.txt (work in progress), Internet Engineering Task Force, January 16, 2006.

[126] H. Schulzrinne, J. Morris, H. Tschofenig, J. Cuellar, J. Polk, and J. Rosenberg. Common policy: An XML document format for expressing privacy preferences. Geopriv Internet-Draft draft-ietf-geopriv-common-policy-10.txt (work in progress), Internet Engineering Task Force, May 21, 2006.

[127] H. Schulzrinne and H. Tschofenig. Location types registry. Geopriv Internet-Draft draft-ietf-geopriv-location-types-registry-06.txt (work in progress), Internet Engineering Task Force, May 21, 2006.

[128] H. Schulzrinne, H. Tschofenig, J. Morris, J. Cuellar, and J. Polk. A document format for expressing privacy preferences for location information. Geopriv Internet-Draft draft-ietf-geopriv-policy-08.txt (work in progress), Internet Engineering Task Force, February 11, 2006.

[129] C. Schwingenschogl and T. Kosch. Geocast enhancements of AODV for vehicular networks. *ACM SIGMOBILE Mobile Computing and Communications Review*, 6(3):96–97, July 2002.

[130] R. Shankaran, V. Varadharajan, and M. Hitchens. Secure distributed location management scheme for mobile hosts. In *Proceedings of the 26th Annual IEEE Conference on Local Computer Networks (LCN 2001)*, pages 296–305. IEEE Computer Society Press, Los Alamitos, CA, USA, November 2001.

[131] A. Smailagic and D. Kogan. Location sensing and privacy in a context-aware computing environment. *IEEE Wireless Communications*, 9(5):10–17, October 2002.

[132] E. Snekkenes. Concepts for personal location privacy policies. In *Proceedings of the 3rd ACM conference on Electronic Commerce (EC '01)*, pages 48–57. ACM Press, New York, NY, USA, October 2001.

[133] M. Spreitzer and M. Theimer. Architectural considerations for scalable, secure, mobile computing with location lnformation. In *Proceedings of the 14th International Conference on Distributed Computing Systems*, pages 29–38. IEEE Computer Society Press, Los Alamitos, CA, USA Press, June 1994.

[134] V. Stanford. Pervasive computing goes the last hundred feet with RFID systems. *IEEE Pervasive Computing*, 2(2):9–14, 2003.

[135] W. Richard Stevens. *TCP/IP Illustrated, Volume 1.* Addison-Wesley Proffessional Computing Series. Addison-Wesley, Indianapolis, IN, USA, 1994.

[136] H. Sugano, S. Fujimoto, G. Klyne, A. Bateman, W. Carr, and J. Peterson. Presence information data format (pidf). RFC 3863, Internet Engineering Task Force, September 2004.

[137] P. Tao, A. Rudys, A. Ladd, and D. S. Wallach. Wireless LAN location-sensing for security applications. In *Proceedings of the 2003 ACM Workshop on Wireless Security (WiSe '03)*, pages 11–20. ACM Press, New York, NY, USA, September 2003.

[138] M. Thomson and J. Winterbottom. Revised civic location format for pidf-lo. Internet-Draft draft-ietf-geopriv-revised-civic-lo-02.txt, Internet Engineering Task Force, April 28, 2006.

[139] Tomtom International BV, http://www.tomtom.com. *Portable GPS car navigation Systems*, April 2006.

[140] Trusted Computing Group. *TPM Main: Part 1 design principles*, 1.2 edition, March 2006.

[141] Trusted Computing Group. *TPM Main: Part 2 TPM Structures*, 1.2 edition, March 2006.

[142] Trusted Computing Group. *TPM Main: Part 3 Commands*, 1.2 edition, March 2006.

[143] H. Tschofenig, F. Adrangi, M. Jones, and A. Lior. Carrying location objects in RADIUS. Geopriv Internet-Draft draft-ietf-geopriv-radius-lo-06.txt (work in progress), Internet Engineering Task Force, March 6, 2006.

[144] Y. Tseng, S. Wu, W. Laio, and C. Chao. Location awareness in ad hoc wireless mobile networks. *Computer*, 34(6):46–52, June 2001.

[145] J.D. Tygar and B.S. Yee. Dyad: A system for using physically secure coprocessors. Technical Report CMU-CS-91-140R, Carnegie Mellon University, May 1991.

[146] US Department of Defense. *Global Positioning System Standard Positioning Service Signal Specification.* US Department of Defense, 2nd edition, June 1995.

[147] V. Varadharajan. Trustworthy computing (extended abstract). In X. Zhou, S. Su, M. P. Papazoglou, M. E. Orlowska, and K. G. Jeffery, editors, *Proceedings of the 5th International Conference on Web Information Systems Engineering (WISE 2004) Brisbane, Australia, November 22-24, 2004*, volume 3306 of *Lecture Notes in Computer Science*, pages 13–16. Springer-Verlag, Berlin, Germany, November 2004.

[148] U. Varshney. Location management support for mobile commerce applications. In *Proceedings of the 1st International Workshop on Mobile Commerce (WMC'01)*, pages 1–10. ACM Press, New York, NY, USA, September 2003.

[149] R. Want, A. Hopper, V. Falcao, and J. Gibbons. The active badge location system. *ACM Transactions on Information Systems*, 10(1):91–102, 1992.

[150] J. Winterbottom, M. Thomson, and H. Tschofenig. GEOPRIV PIDF-LO usage clarification, considerations and recommendations. Geopriv Internet-Draft

draft-ietf-geopriv-pdif-lo-profile-04.txt (work in progress), Internet Engineering Task Force, May 2, 2006.

[151] P. Yau and C. J. Mitchell. 2HARP: A secure routing protocol to detect failed and selfish nodes in mobile ad hoc networks. In *Proceedings of the 5th World Wireless Congress*, pages 1–6. Delson Group Inc., San Francisco, CA, USA, May 2004.

[152] J. Zagami, S. A. Parl, J. Bussgang, and K. D. Melillo. Providing universal locations services using a wireless E911 location network. *IEEE Communications Magazine*, 36(4):66–71, April 1998.

[153] X. Zhang, F. Parisi-Presicce, R. Sandhu, and J. Park. Formal model and policy specification of usage control. *ACM Transactions on Information and System Security (TISSEC)*, 8:351–387, 2005.

[154] F. Zhu, M. W. Mutka, and L. M. Ni. Splendor: A secure, private, and location-aware service discovery protocol supporting mobile services. In *Proceedings of the First IEEE International Conference on Pervasive Computing and Communications*, pages 235–242. IEEE Computer Society Press, Los Alamitos, CA, USA, March 2003.

Appendix A

Information from web-pages

Because the policies of service providers are continually changing, and also because of the difficulties in finding such information after it has changed, relevant extracts from web pages have been reproduced here to provide a consistent basis for this thesis. The information in the Appendix is referred to in Section 3.5.2.

A.1 Introduction

The text in Appendix A.2 is an extract from *http://www.cingular.com/privacy/privacy_policy*. The text in Appendix A.3 is an extract from *http://www.cingular.com/mmode/mmode_net*. Both texts were copied on 2 September 2005.

A.2 Cingular Wireless Privacy Policy

Cingular Wireless uses e-mail, short text messages, telemarketing, and direct mail to inform you about products or services we think will interest you. If you do not wish to receive these types of communications, learn how you can opt out. (Effective date: September 2, 2005)

Cingular Wireless has a long-standing policy of protecting customer privacy.

We believe that you should know what information we collect from you, as well as how that information is used, disclosed, and protected. We have created this policy

statement (the "Policy") to explain our privacy practices and policies.

We will not sell or disclose your personal information to unaffiliated third parties without your consent except as otherwise provided in this Policy. We may use information about who you are, where and when you browse on the Web, where your wireless device is located, and how you use our network to provide you better service and enrich your user experience when you sign up or use any of our products or services.

Cingular Wireless may make available shorter or machine-readable versions of this Privacy Policy. These additional policies are intended only as summaries of this complete Privacy Policy. This Policy applies to customers who purchase and use our services and products in the United States. Collection, use, disclosure, and protection of personal information may be subject to different regulation outside the United States.

Cingular Wireless will revise and update this Policy as it deems appropriate, including, for example, if our practices change or if changes in the law so require. You should refer back to this page for the latest information.

QUICK PRIVACY LINKS

The following brief summaries outline our Policy and then link to further detail. We also provide you the means to communicate with us if you have any questions about this Policy. This Policy also addresses what we do with information about your device usage, the services you buy from us, and who you call.

Does Cingular Wireless Collect Personal Information About Me?

Yes, Cingular Wireless collects personal information about you so that we can deliver products or services you request. This happens automatically when you interact with us, such as when you log into a service. Sometimes, Cingular Wireless buys commercially available marketing and sales information from third parties so we can

better serve you.

Does Cingular Wireless Collect Information About Children Under the Age of 13?

If we make available offers and products online where a child informs us that he or she is under the age of 13, we will ask a parent to confirm his/her consent in advance of further collection, use or disclosure of personal information from that child.

How Does Cingular Wireless Use Personal Information?

We use personal information for billing and collection purposes, to provide services or complete transactions you have requested, and to anticipate and resolve problems with your services. We may also use this information to create and inform you of products or services from Cingular Wireless or others that may better meet your needs.

When Does Cingular Wireless Disclose Personal Information?

We do not sell personal information to unaffiliated third parties. We will disclose personal information to third parties to complete a transaction you have requested, as part of the terms and conditions for a particular service, to collect on an account, or when we otherwise have your consent to do so. We also may disclose personal information to third parties to protect the rights and property of the company or its subscribers.

What Happens to Information About My Telephone Usage and Who I Call?

Under federal law, you have a right, and we have a duty, to protect the confidentiality of information about your device usage, the services you buy from us, who you call, and the location of your device on our network when you make a voice call.

Does Cingular Wireless Use Cookies?

Yes. Cingular Wireless uses cookies for a variety of reasons, including improving your shopping or browsing experience. In addition, cookies help us personalize the site experience for you.

Does Cingular Wireless Use Web Beacons?

Yes. Cingular Wireless uses Web beacons, also known as Web bugs, on our sites so that we can identify you and deliver you the services you request. However, Cingular Wireless does not permit third parties to use Web beacons linked to personal information on our site.

Does Cingular Wireless Place Advertising on Other Web Sites?

Cingular Wireless currently uses third-party advertising companies to place our ads on the Internet, and cookie and Web beacon technologies are used to measure the effectiveness of those ads. You should know that the use of such cookies is subject to the third parties' privacy policies, and not the policy of Cingular Wireless.

What About Presence, Location, and Tracking Information?

Our network knows the general location of your phone whenever it is turned on. When we offer you optional services that require use or disclosure of this information, the terms and conditions for the specific service offering explain how the location information will be used. We also may provide your network location to emergency service providers if you place a 911 call.

Can I Choose Not to Receive Marketing Messages on My Wireless Device from Third Parties?

Cingular Wireless has implemented technology to reduce unsolicited bulk short text messages but is unable to filter all marketing messages that you receive on your wireless device from third parties.

How Secure Is Information About Me?

We maintain a variety of physical, electronic, and procedural safeguards to guard your personal information.

What Can I Do to Protect My Personal Information?

An important part of ensuring the security of your personal information is your own efforts to protect against unauthorized access to your wireless device and SIM card. Before discarding your device or trading it in, be sure you remove all your personal information from the device.

How can I Review Personal Information in my Account for Accuracy?

You can review the accuracy of the personal information in your account online (go to My Account on this Web site) or by contacting Customer Service at 1-866-CINGULAR.

Will This Policy Be Updated?

Cingular Wireless expects to update this Policy periodically. You should refer back to this page often for the latest information and the effective date of any changes to the Policy.

To Whom Should I Direct Privacy Questions or Concerns?

You have several ways to contact us about questions to this Policy and about your services.

CINGULAR WIRELESS PRIVACY POLICY

We have created the Cingular Wireless Privacy Policy (the "Policy") to explain our privacy practices. When you use any Cingular Wireless product or service, you should understand when and how personal information is collected, used, disclosed and protected.

We will not sell or disclose information to unaffiliated third parties without your consent except as otherwise provided in this Policy. We may use information about who you are, where and when you browse on the Web, where your wireless device is located, and how you use our network to provide you better service and enrich your user experience when you sign up or use any of our products or services.

Cingular Wireless will revise and update this Policy if our practices change or if changes in the law so require.

Information Collected About You

We collect a variety of personal information about users of our products or services. Personal information is information that can be directly associated with a specific person or entity, such as a name, address, telephone number, e-mail address, or information about activities directly linked to that person.

Our definition of personal information does not include "aggregate" information. Aggregate information is data we collect about a group or category of services or customers from which individual customer identities have been removed. For example, we could prepare a report that indicates that a certain number of our customers always use their wireless phones at a certain time of day at a specific location. Ag-

gregate data helps us understand trends and customer needs so that we can better consider new services or tailor existing services to customer desires. The aggregate data also might be purchased by or shared with a third party, for example, one interested in locating a business in a particular part of town.

Here are the types of personal and other information we collect. You should refer to the rest of this Policy to see how we use, disclose, and protect that information:

- **Information You Give Us:** We collect information you give us when you purchase a Cingular Wireless product or use services. For example, you may provide us a billing address and credit information, including your social security number or business identifier, when signing up for service or perhaps purchasing a product through our online store. You might not have thought about it this way, but the numbers dialed from your wireless phone to make a call are an example of information you give us and that we collect and use so we can bill you appropriately and investigate fraudulent usage.
- **Automatically Collected Information:** We automatically receive certain types of information whenever you interact with us. For example, when you visit a Cingular Wireless Web site, our systems automatically collect your IP address and the type of browser you use. When you browse the wireless web, our systems log the Web sites you visit. Similarly, all wireless communications systems know when your phone is turned on and approximately where the device is physically locatedthat's how calls or messages are delivered to you in real time.
- **Information from Other Sources:** We may obtain information about you from outside sources and add it to or combine it with your account information. For example, we may receive credit information for purposes of initiating service. We also may use commercially available demographic and marketing information from third parties to help us better serve you or inform you about products or services that we think will be of interest to you. We sometimes receive updated delivery and address information from our shippers or other sources so that we can correct our records and deliver your next purchase or communication more easily. And, we often receive information from the dealer from whom you purchase your wireless phone or device prior to initiating ser-

vice with us.

We also may purchase e-mail lists from third parties for advertising purposes. We only purchase lists of individuals who have allowed third-party use of their e-mail address for marketing purposes. If you have previously requested to participate in an e-mail advertising program, the information we receive may include your name, information on previous transactions, or any other personal information you have provided.

Cingular Wireless recognizes that parents often purchase our products and services for family use, including for use by minors. Any information collected from such usage will appear to be the personal information of the actual subscriber to the service, and will be treated as such under this Policy. If we make available offers and products online where a child informs us that he or she is under the age of 13, we will ask a parent to confirm his/her consent in advance of further collection, use or disclosure of personal information from that child.

In a Business Agreement, our customer is a business or other entity purchasing service for employees or other authorized users. If you receive certain benefits through a business or government customer's agreement with us, this Policy will generally govern your personal information. However, if you receive service where a business or government entity pays Cingular for your account or is otherwise liable to Cingular for the charges (for example, as a guarantor if you fail to pay), we may share your account information with that entity. If you receive certain benefits tied to a Business Agreement, but you are liable for your own charges, then we may share enough account information with that entity to verify your continuing eligibility for those benefits. Please contact Cingular if you have any questions about who is the liable party on your bill.

Children's Online Privacy Protection Act

If we make available offers and products online where a child informs us that he or she is under the age of 13, we will ask a parent to confirm his/her consent in advance of further collection, use or disclosure of personal information from that child. You

should be aware, however, that wireless devices and services purchased for family use may be used by minors without the knowledge of Cingular Wireless. If that happens, any information collected from the usage will appear to be the personal information of the actual adult subscriber and treated as such under this Policy.

Use of Personal Information

- **Internal Use.** In general, we use personal information to serve our customers, to enhance and extend our customer relationship, and to enable our customers to take maximum advantage of products and services we think they would enjoy. For example, by understanding how you use our Web site, we are able to customize and personalize your experience. More specifically, we use personal information for billing purposes, to provide services or complete transactions you have requested, to anticipate and resolve problems with your services, and to create and inform you of products or services from Cingular Wireless or others that better meet your needs.

  Cingular Wireless uses e-mail, short text messages, telemarketing, and direct mail to inform you about products or services we think will interest you. You can modify your preferences on receiving these types of communications:

  - If you were formerly an AT&T Wireless customer and wish to change your preferences on receiving these types of communications, you can complete an online form to tell us your preferences.
  - Other Cingular customers who wish to change their preferences can visit the "My Profile" page, if you manage your account online, or call Customer Service at 1-866-CINGULAR. If you are not a current customer of Cingular Wireless and wish to opt out of receiving marketing communications, you can complete an online form to tell us your preferences.

  While you may choose not to receive marketing information from us, you will continue to receive invoices, customer service-related notifications, and other similar information from us electronically or otherwise.

- **Third-Party Use.** You should review the following section to understand when Cingular Wireless discloses personal information to third parties.

Disclosure of Personal Information

Information about our customers is one of our most important business assets, and therefore we strive to protect it and keep it confidential. We do not sell personal information to third parties without your consent. When and what types of information Cingular Wireless disclose depend on the service and in some cases the choices you have made.

Cingular Wireless will not disclose personal information other than in accordance with this Policy. In general, that means that you must consent to the disclosure in advance. Depending on the service, we may obtain your consent in a number of ways, including:

- In writing;
- Verbally;
- Online by clicking a button;
- Through the use of a dialing string or button on a wireless device or handset; or
- At the time of initiation of a particular service offering, when your consent is part of the required terms and conditions to use that service.

For example, your consent to disclose personal information can be implied simply by the nature of your request, such as when you ask us to deliver an e-mail or short message to another person. Your return address is disclosed as part of the service and your consent to do so implied by your use of the service. To determine how personal information may be disclosed as part of a particular service, you should review the terms and conditions of use for that service.

We share personal information with third parties as necessary to complete a transaction, perform a service on our behalf (such as enhancing our ability to serve you better), or perform a service that you have requested. When the third party acts solely on our behalf, Cingular Wireless does not allow them to use your information

for other purposes. For example, our vendors process and print your billing statement on our behalf. They can only use the personal information we give them to produce the billing statement. When we write off an account for non-payment, Cingular Wireless sometimes disclose personal information about the account to third parties such as credit bureaus. Credit bureaus may use the personal information to update their records. When we write off an account for non-payment, Cingular Wireless sometimes discloses personal information about the account to third parties such as credit bureaus. Credit bureaus may use the personal information to update their records. Cingular Wireless does not currently disclose wireless numbers in directory assistance listings or published directories. If we do so in the future, you will be able to choose whether your number is listed.

Aside from our services, however, you may also want to take advantage of services and products offered by other companies utilizing our wireless service. In those cases, you will be providing information to those companies, and information about you received by those third parties will be governed by their privacy policies, not this Policy. For example, if you are roaming on the network of another carrier, information about your usage and the numbers you dial will be available to the carrier providing the service. Also, as another example, if you purchase something using our mobile Internet service, you will be disclosing personal information directly to the company facilitating the transaction, a merchant bank and the merchant. Finally, if you bought your wireless device from a third party retailer or dealer, both they and Cingular Wireless will have personal information as a result of the transaction and your ongoing service with Cingular. Whenever third parties have a role in any such transaction, you should review their privacy policies as well.

From time to time you may be able to participate in contests, giveaways, or other similar promotions we sponsor. Except as explained otherwise in the rules for a particular contest, giveaway, or promotion, any personal information you provide will be used in accordance with this Policy.

In addition, from time to time you may be able to participate in our surveys to help us improve our offerings and services. Except as explained otherwise in the survey, any personal information so collected will be used for our internal purposes.

Under federal law, you have a right, and we have a duty, to protect the confidentiality of information about your telephone usage, the services you buy from us, who you call, and the location of your device on our network when you make a voice call. This information is sometimes referred to as "Customer Proprietary Network Information," or "CPNI." We share CPNI and other personal information about you with affiliates of SBC Communications and BellSouth Corporation (the parent companies of Cingular) that provide telecommunications services to which you also subscribe. Before sharing CPNI in any other way, we will first notify you of your rights under the law, describe how we intend to use the CPNI, and give you an opportunity to opt out of such usage (or, when required by law, to opt in).

e-Wallet Services:

Customers who use our e-Wallet services to purchase products or services with their credit card, debit card, or have the charges placed on their wireless bill, should review our e-Wallet Supplemental Privacy Notice to learn about our policies and practices for the treatment of personal information about you that is collected through your use of the e-Wallet services.

Business Transfers:

Information about our users, including personal information, may be disclosed as part of any merger, acquisition, sale of company assets, or transition of service to another provider, as well as in the unlikely event of an insolvency, bankruptcy, or receivership in which personal information would be transferred as one of the business assets of the company.

Protection of Cingular Wireless and Others:

We release personal information when we believe release is appropriate to comply with the law or in good faith reliance on legal process (e.g., court orders, subpoenas, E911 information, etc.); enforce or apply our customer agreements; initiate, render,

bill, and collect for services; protect our rights or property, or protect users of those services and other carriers from fraudulent, abusive, or unlawful use of, or subscription to, such services; facilitate or verify the appropriate calculation of taxes, fees, or other obligations due to a local, state, or federal government; or if we reasonably believe that an emergency involving immediate danger of death or serious physical injury to any person requires disclosure of communications, or justifies disclosure of records, to a governmental entity without delay.

Cingular Wireless Use of Cookies

Cookies are small files placed on your computer's hard drive by a Web site when you visit. These files identify your computer and record your preferences and other data about your visit so that when you return to the site, the site knows who you are and can personalize your visit. For example, cookies enable a Web site shopping cart to function so that you only have to check out once. Consequently, cookies are often tied to the use of personally identifiable information while using our site, and some functionality may be lost if you choose not to accept the use of cookies.

In general, we use cookies to collect information so that we can determine how to improve our site by seeing which areas, features and products are most popular; to personalize the site and make recommendations based on products you have liked in the past as well as to improve the site experience; and to complete transactions you have requested. Advertisers that serve ads on our site may also use their own cookies. Such outside cookies are governed by the privacy policies of the entities placing the ads, and are not subject to this Policy.

We also use a session-based cookie that maintains a user's session for continuity of navigation while viewing the site. After closing the browser the session cookie simply terminates.

With wireless Internet service, cookies are also used by our suppliers and third-party vendors to facilitate the various services and information offered. Depending on the phone or device being used, cookies may be stored locally on the phone or device, or on servers operated by Cingular Wireless. This Internet cookie-like functionality is

in place for the same reasons and designed for the same purposes as cookies placed on your computer when interacting with Internet Web sites. Our suppliers and vendors will only use this information to provision the service, but each site you visit while using the wireless Internet service is controlled by a separate company and their individual privacy policies will govern information they receive automatically from the cookie or information you voluntarily provide.

Cingular Wireless Use of Web Beacons

A Web beacon, also known as a Web bug, is a small, graphic image on a Web page, Web-based document or in an e-mail message that is designed to allow the site owner or a third party to monitor the address and other information of the computer viewing the item. Web beacons are often invisible to the user because they are typically very small (only 1-by-1 pixel) and the same color as the background of the Web page, document, or e-mail message. Web beacons are represented as HTML IMG tags in the Web page; users can click on "view profiles" of the Web page to see whether the page is using a Web beacon. Web beacons collect the IP address of the computer that the Web beacon is sent to, the URL of the page the Web beacon comes from, and the time it was viewed. Web beacons can also be linked to personal information.

Cingular Wireless does not place Web beacons that link to personal information on other sites, nor does it permit third parties, other than those working on our behalf, to place them on our site. Cingular Wireless does use Web beacons itself and may link a particular beacon to personal information. For example, we may use a beacon to ensure a user can flip between technical assistance, customer service and our online store and still be recognized as our customer.

Cingular Wireless Placement of Advertising on Other Web Sites

Cingular Wireless may use third-party ad serving companies to place advertisements about our products and services on other Web sites. These companies may use cookies and other technology such as Web beacons or tagging to measure the effectiveness

of our ads. To measure advertising effectiveness and offer selective ad content, the ad serving companies may use anonymous information about your visits to our and other Web sites. But the ad serving companies use an anonymous number to identify you, NOT your name, address, phone number, e-mail address, or anything that personally identifies you. The use of such cookies is subject to the ad serving company's privacy policy, not the Policy of Cingular Wireless. If you would like more information about these companies we use, their privacy practices, or to learn your choices about not having this non-personal information used to serve ads to you, see our Cingular Wireless Cingular Wireless Ad Serving Companies.

Presence, Location, and Tracking

To make wireless communications possible, the network knows the general location of your phone or wireless device whenever it is turned on. Your wireless device sends out a periodic signal to the nearest radio tower/cell site so that the network will know where to route an incoming communication and how to properly bill for the service. This is necessary to make wireless communications possible.

If you dial 911 for emergency services, we may provide your network location to a public safety answering point, emergency medical service provider or emergency dispatch provider, public safety, fire service, or law enforcement official, or hospital emergency or trauma care facility. The law also permits us to disclose the location of a device on our network without a user's consent (1) to a user's legal guardian or members of a user's immediate family in an emergency situation that involves the risk of death or serious physical harm, (2) to database management services or information providers solely to assist in delivering emergency services, or (3) to a governmental entity if we reasonably believe that an emergency involving immediate danger of death or serious physical injury to any person requires or justifies disclosure of a device's location on the network without delay. Legally required upgrades will allow us to provide a location more precise than cell site location.

In addition, we offer optional services on our GSM/GPRS network that make use of your network location. Please review the terms and conditions for each service for additional information about how the location information will be used. The location

used for these services is separate from the network location information when you make a voice call. Your wireless Internet service may also be personalized using your ZIP code or other location identifiers. We use this information to serve you relevant content, and we treat the information like any other personal information under this Policy. This service does not use the network location technology described in this section.

Receipt of Marketing Messages on My Wireless Device from Third Parties

You should be aware that not all advertisements appearing on or delivered to your mobile phone or device are authorized by Cingular Wireless. We have developed and implemented systems in our network to reduce unsolicited bulk short text messages, but we cannot at this time block all such messages. We continue to look for other options to reduce these unsolicited bulk messages. If you have an e-mail account with Cingular Wireless, this service is subject to unsolicited messages as any other e-mail service.

It is unlawful for any third party to make an unsolicited telemarketing call using an autodialer or to send a prerecorded message to a wireless phone or device. You should report any such unsolicited calls to the Federal Communications Commission.

Network and Information Security

We maintain a variety of physical, electronic, and procedural safeguards to guard your personal information. For example, we use accepted tools and techniques to protect against unauthorized access to our systems. Also, we grant access to personal information about you to employees and contractors who need to know that information to provide products or services to you. In addition, we work to protect the security of your personal information when you are ordering new service via the Cingular Wireless Web site by using well-known Internet encryption technologies like Secure Sockets Layer (SSL). We also use encryption technologies to protect your account information when you are viewing your bill on our Web site. You should be aware that Cingular Wireless has no control over the security of other sites on the

Internet you might visit, interact with, or from which you buy products or services.

What Can I Do to Protect My Personal Information?

An important part of ensuring the security of personal information is your own effort to protect against unauthorized access to your wireless device and the personal information contained in it and on your SIM card. Most phones and wireless PDA-type devices store calling information both in the phone and on the SIM card. Therefore, before discarding your phone or PDA, trading it in or giving it away, be sure you remove and retain your SIM card and follow the manufacturer's instructions for deleting all personal information on the device itself. (This can be found in your owner's manual or on the manufacturers' Web site.)

In addition, use passwords to prevent unauthorized access to your wireless device, your wireless service account, and your voicemail. If you write down your passwords or user names, keep the information in a secure location. Do not give your password to someone else unless you intend them to have the same full access and ability to make changes to your account as you have. Change your passwords periodically.

Accuracy of Personal Information in Your Account

You can review the accuracy of the personal information in your account records online (go to My Account on this Web site) or by contacting Customer Service at 1-866-CINGULAR.

California Customers

Cingular Wireless does not disclose customers' personal information to third parties for the third parties' direct marketing purposes, as governed by California Civil Code 1798.83.

Updating this Policy

Cingular Wireless will revise or update this Policy as it deems appropriate, including for example, if our practices change, as we change existing or add new services, as we develop better ways to inform you of products we think will be of interest, or if the law so requires. You should refer back to this page often for the latest information and the effective date of any changes.

If, however, users' personally identifiable information will be used in a manner materially different from that stated at the time of collection, we will notify users via posting on this page for 30 days before the material change is made. Users will have a choice as to whether or not their information will be used in this materially different manner.

Contact Us

Cingular Wireless is committed to the policies set forth in this Policy.

If you have any questions, comments or concerns about this Policy, please contact privacy@cingular.com.

If you have questions about your Cingular Wireless service, you can call a customer service representative at 1-866-CINGULAR.

A.3 Features and Services Information for Former AT&T Wireless Users - mMode

Features & Benefits

mMode brings the wireless Internet to your phoneallowing you to download ringtones and games, check e-mail, and stay informed. Content providers span the spectrum from news, finance, and sports, to weather, traffic, and shopping. With mMode,

there's something for everyone.

E-mail with mMode

Use your phone to send and receive e-mail messages from AOL, MSN Hotmail, Yahoo! Mail, and more.

Getting mMode

Whether you'll use mMode all the time, just when you're on the go, or only for emergencies, there is an mMode plan to suit your needs. When you use mMode to send e-mail, read news, or download a game, your wireless phone works like a computer to transmit and receive data.

View our mMode plans to determine the plan that fits your needs. To set up your mMode account, call 1-877-400-1080.

My mMode

Personalize your mMode experience with the sites you use most. The world of mMode is vast, so having your own personalized world simplifies your needs. That's My mModeeasy access to your account, your way. Create it. Change it. My mMode shifts as your life does.