Location Privacy in the Alipes Platform

Kåre Synnes, James Nord, Peter Parnes
Luleå University of Technology / Centre for Distance-spanning Technology
Department of Computer Science and Electrical Engineering
SE-971 87 Luleå, Sweden.
{Kare.Synnes, James.Nord, Peter.Parnes}@cdt.luth.se

Abstract

An increasing number systems use contextual information about their users. Such contextual information can be used to design applications that survey usage and adapt thereafter, or simply just use context information to optimize presentation. Context information could therefore be used to create applications for the benefit of the users of the system, but the same information could cause serious violations of personal integrity if misused. Locality may be the most widely used, but also the most sensitive contextual information. The Alipes platform makes it easy to create location-based services while enforcing user privacy and integrity. The platform handles privacy through rules that describe how and under what circumstances a user's context may be distributed to other users, for example rules describing limitations concerning the user's context, a certain time period, the number of queries and the type of applications. This paper presents how location privacy is enforced in the Alipes platform.

1. Introduction

An intriguing question is how mobile applications differ from other applications, mobile applications being defined as applications used in a mobile context such as in a mobile terminal. The simple answer is that the difference is minimal, if there is a difference at all, since a user would like to do anything that he could do at a stationary terminal from his mobile terminal. However, from another point of view the difference can be said to be large. The applications may be similar or even identical if used in mobile or stationary terminals, but the usage may differ depending on the context of the user.

For instance, using a stationary computer at work basically determines the main context of usage of that terminal as being work, while using a stationary terminal from home would instead imply another main context, namely leisure. If the user also has a mobile terminal then that terminal may switch its main context of usage, between work and leisure, depending on the current location. The situation is more complex in real life than this example may demonstrate, as many users spend time on leisure activities while at work and some users work from home, so the current location is but one factor that may help determine the main context of a user.

What real benefits can come from using a user's context, such as his location, in mobile terminals? It has some obvious benefits, such as enabling traditional map services where the user can locate himself or relevant nearby services like the closest bus stop, a friend or a restaurant. Another benefit may be to optimize the mobile terminal to use the most suitable set of tools for a particular context. An additional benefit of the user's location could be derived in a protective system, where the user could hit an alarm when assaulted so that help could be sent to his location, or using automatic alarms in combination with sensory technology for determining glucose levels, heart-beat rates, etc. Finally, it may benefit computer gamers of all ages, by creating mobile games based on the user's physical location and context.

The user's location and other context information can also be used for less admirable purposes. Someone could, for instance, trace the habits of a person, such as his movement patterns, and thereby that deduce information that is highly personal and thus private. The classic examples are detecting an unfaithful spouse or monitoring workers. It is therefore vital to consider privacy issues when designing distributed real-time systems where context information may be distributed to a number of users. Privacy can be defined as the demands from of individuals, groups and institutions to determine by themselves when, how and to what extent information about them is to be communicated to others [1]. Köhntopp et al define related terminology such as anonymity, unlinkability, unobservability, and pseudonymity [2]. It is our opinion that a system can be seen as a threat to users' integrity if privacy cannot be enforced. Such systems will not be trusted and will therefore probably not be used by the majority of users.

Ackerman et al state, "Indeed privacy forms a co-design space between the social, the technical, and the regulatory" [3]. It is therefore also important to consider national, regional and international legislative initiatives and interest groups. The Swedish Data Protection Acts [4-

5] needs to be considered when applying systems in Sweden. Legislative initiatives to control sensitive information, such as location data, are also under way both in the United States and within the European Union [6-10]. In short, they specify that users of mobile location services must be protected by privacy safeguards, must be fully informed of the purposes of the usage of the mobile location services, and must have the right to determine the use of their personal information. The European telecommunication directives discuss the idea of a user's right to choose whether to grant (opt-in) or deny (opt-out) any use of information regarding himself at any time. This also leads to the possibility of a user temporary disabling a service, if he wants to do so.

Privacy has also been discussed by the Location Interoperability Forum [11], which is a joint effort initiated by telecom vendors, service providers and operators to create a common location interchange format, but little of that work is as yet public or freely available. The same is true of other industry initiatives like the Location Positioning Workgroup or the WAP Forum [12], whose proceedings are unfortunately confidential. There is also the IETF Geographical Location/Privacy Workgroup (GeoPriv) [13-14], which is working on standards and recommendations for location-based services.

When studying directives and proposals regarding privacy within Sweden and the European Union, it becomes clear that a positioning platform, such as the Alipes architecture [15], must also take personal integrity issues into account. This is especially noteworthy when one also considers the numerous international initiatives and human rights groups [16-20], as they symbolize the deficit of current legislation as well as the mistrust of governments and legislative bodies. For example, Westin found that 81% of Net users are generally concerned about threats to their privacy while online [21].

It is clear that the user should have control over his information, and therefore the ownership of the information itself should belong to the user. Anyone who wishes to use that information can be given rights, but the user is the only logical owner of his own private information.

Studies have shown that users can be divided into three major clusters with regard to privacy: the privacy fundamentalists (17%), the pragmatic majority (56%), and the marginally concerned (27%) [22-23]. These clusters view privacy differently, which has to be taken into consideration when creating a mechanism for enforcing privacy.

Section 2 describes the Alipes platform for seamless interchange of position information, while Section 3 discusses privacy enforcements by the platform. In Section 4 some of the topics for future work are presented and Section 5 concludes the paper with a summary.

2. The Alipes Platform

The Alipes platform is an architecture for creating applications aware of the user's context, where the location is the most important factor [14]. The platform is mainly intended for real-time use, i.e. for applications that require a constant update of the user's context. It is also inherently distributed, meaning that clients can communicate directly with each other without a central server being present.

There are several advantages of building distributed systems, as a central server may become a bottleneck as the usage of a system increases. A better alternative is to design the clients of the system to be autonomous and thereby create a system that may scale more easily. However, some situations, such as locating a user the first time, may however require a central server or a topology of servers, as described further down in this section.

The following subsections describe the Alipes platform in more detail in order to give a better understanding of the privacy implications. It is especially important to understand the process of how a user is positioned and how location servers can be used together with the platform. Section 2.2 may be read as an overview and is included here for completeness.

2.1. The Alipes Platform

Mobile terminals have recently become both networked and powerful enough to host distributed real-time applications such as video-conferencing tools [24]. As argued in the introduction, the perhaps only thing that distinguishes a mobile terminal from a stationary one is its use in several contexts, since it is kept with the user while for instance at work or at home.

This is important for two reasons. First, the context can be used to optimize the view on the mobile terminal to suit the current use of the terminal; at work information and applications are filtered out to present the most efficient environment for work, while at home quite a different presentation may be the most efficient from the perspective of the user's interests and entertainment. This could help overcome the increased effects of stress and information overload that today plague modern society. Secondly and more importantly, some context information can be used to create a new range of applications that depend, for instance, on the location of the user. These location-aware applications span from simple 'where-am-I' applications that include map services and 'friend-finder' tools, electronic guides and mobile learning systems, to systems including other sensory technologies, which can, for instance, automatically send a call for help when a user is in a hazardous situation (such as having a heart-attack or has simply have fallen down). Another area of interest is the growing area of computer gaming, where games like "Pirates!" are based on location information [25].

Most applications on the market are written for a specific positioning technology such as GPS, but given the flexibility of modern mobile terminals, other positioning technologies could also be used. The Alipes platform allow multiple positioning techniques to be seamlessly interchanged and combined, enabling applications to utilize a single interface, yet benefit from advantages that no single positioning technique can offer by itself. Several positioning techniques could thus be combined to achieve benefits such as full indoor and outdoor coverage, or to determine more accurate positions.

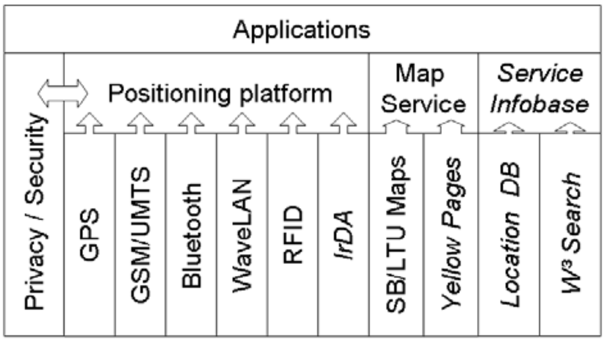

Figure 1. The Alipes platform.

The platform has currently four positioning techniques implemented, namely GPS, GSM/UMTS (MPS), Bluetooth and WaveLAN, while work on adding Radio Frequency Identification (RFID) has been started and work on infrared (IrDA) is planned [15,26-27]. Figure 1 above depicts the platform with the as yet unimplemented modules in Italics. The platform is designed to use different map services based on OpenGIS or other standards [28]. The figure also shows that different service information databases can be used to retrieve information about objects, such as location servers containing information on objects and their related location. How the privacy and security aspects fit into the platform will be described in Section 3.

The platform also allows peer-to-peer interchange of position information using ad-hoc networks and thus offers a wide variety of techniques to be interchanged or combined, with obvious advantages such as possibly greater accuracy and better coverage. Figure 2 shows two users sharing position information to achieve higher precision (ε is the range of the Bluetooth device), for example Carol merging her MPS information with Dave's WaveLAN information while considering the error of the added peer network range. The interchange is achieved without exchanging identities, thus allowing anonymity.

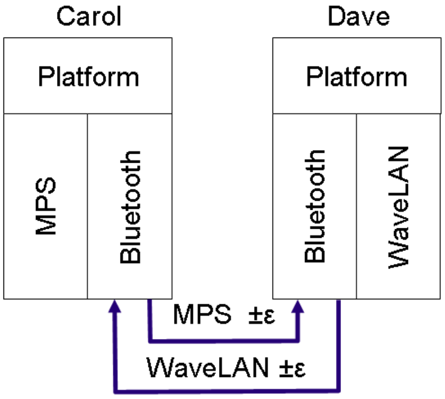

Figure 2. Peer-to-peer position exchange using ad-hoc networking.

The platform is mainly designed for distributed clients, and no or few servers need to be involved in the system. Each client relies on peer-to-peer networking to exchange information, but the clients also need networks to retrieve data from databases (such as map servers and location information databases). The clients could thus be designed to be fully symmetric and distributed in such a way that they are the intelligent nodes and the network is more or less dumb.

2.2. Locating a User

The user can locate himself using the platform, as described in Section 2.1, but he might also want to locate other persons. An example could be a 'friend-finder' application, where his close friends' locations and context are presented on a map. Using that information together with a conferencing or messaging system (see [24]) would enable him to find which of his friends are free for lunch, for instance. Locating a user would involve a number of consecutive steps:

1. Finding the ID of the user: The ID typically consists of a username and a terminal or domain name, such as *unicorn@porthos.cdt.luth.se*. In this case *unicorn* is the username, alias, or pseudonym selected by the user, *porthos* is the name of the mobile terminal and *cdt.luth.se* is the domain name of the mobile terminal. The terminal name *porthos.cdt.luth.se* could also be linked to a mobile IP address [29] or to an IP multicast address [30]. Note that the ID could be anything that is unique, but there need to be ways to distinguish between multiple terminals. There are several ways to discover the ID of the user, ranging from using email or chat tools to a directory service or specialized applications.

2. Finding the IP address of the user terminal: The IP address could be found in several ways by using the ID of the user.
    a. If it is based on a mobile IP address, then no further look up is necessary.
    b. If it is based on a static IP address, then DNS can be used to look up the IP address from the hostname , such as from porthos.cdt.luth.se to 130.240.64.72.
    c. If the user is on a local ad-hoc network, then a local broadcast or multicast query could be send out with the TTL (Time-To-Live) set low on a fixed IP multicast or broadcast address. The client would then listen to queries on that address and reply with the related IP address.
    d. A wide-scope IP multicast query could also be used where the TTL is set high and the domain cdt.luth.se is associated with a certain IP multicast address for queries. The client would again listen to queries on that address and reply with the related IP address.
    e. Finally, a look-up of the related IP address can be carried out using central servers or a topology of servers. The Session Initiation Protocol (SIP) [31] or the Lightweight Directory Access Protocol (LDAP) [32] both provide mechanisms for this purpose.
3. Finding the location of the user: When we have the IP address of the user's client, we can then send queries about the user's location. How this is achieved is described in Section 3.

Most of this can be done in a distributed way without including any central server in the schema and thus also avoiding possible bottlenecks. Using IP multicast or mobile IP is elegant solutions, but these technologies are not yet widespread. Consequently, in some cases resolving the IP address of a user's terminal requires using a central server or a topology of servers. SIP and LDAP are inherently based on central resources, even if at least SIP could be said to be somewhat distributed. SIP is perhaps the best scheme, as it is designed exactly for this purpose.

Figure 3 shows an example where Carol is using four positioning techniques through the Alipes platform: GPS, WLAN (WaveLAN), BT (Bluetooth) and MPS (GSM/UMTS). Carol is identified by *carol@cdt.luth.se* and would be located by Dave in three steps when:

1. When locating the responsible server, by using DNS or other mechanisms, we find the SIP server *sip.cdt.luth.se*.
2. Now we can *query sip.cdt.luth.se* for the current IP address of the user, which is *ipaq.homeip.net*.
3. The last step is to communicate directly with the client.

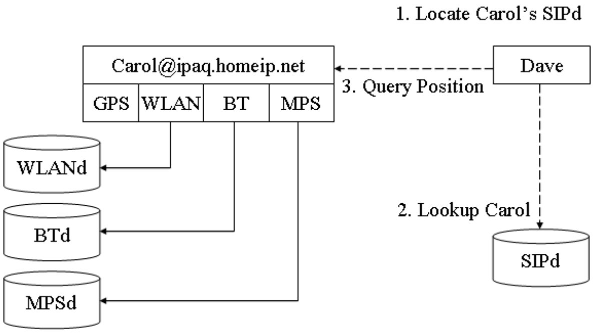

Figure 3. Dave locating Carol using SIP and the Alipes platform.

2.3. Position and Location Servers

A positioning server is normally bound to one positioning technique and offers a positioning service to clients or other services. A positioning service could of course aggregate several other positioning servers and thus offer a wider range of services. Three positioning servers (WLANd, BTd and MPSd) are included in Figure 3 above.

A location server is a server where the location (position) of objects and users is primarily reported by the client terminal itself. A location server may also be setup to use positioning servers in the network directly. The client would then report the position gained from local devices, such as RFID, GPS or Bluetooth devices, to the location server. The location server would report the position gained from the positioning servers in the network to the client.

Figure 4 shows a co-located location server (LOCd) and SIP server, where the location server also handles WLAN, Bluetooth and MPS positioning.

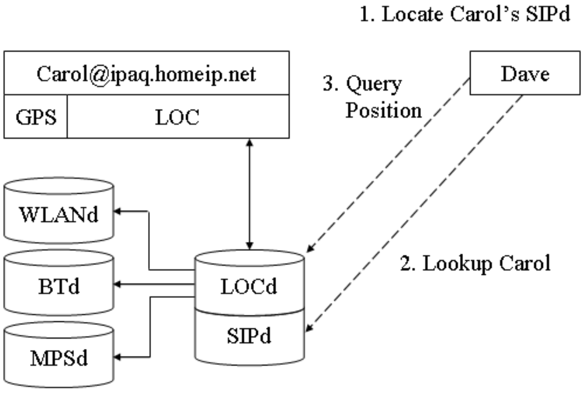

Figure 4. Dave locating Carol using SIP, a location server and the Alipes platform.

The main benefit of a location server is when the client is behind a firewall or is using network address translation that effectively makes queries to the client itself very difficult. The client terminal may also be offline or turned off while the user himself may be possible to position (for example if he has a GSM phone). The client may also be a thin client like a small PDA or have limited network capabilities, which favor using an intermediate entity to limit the mobile terminal's requirement for processing and network usage.

3. Location Privacy

A simple example of the use of contextual information could be presenting information about historical artifacts of interest in the close vicinity of a user, based on the user's location and how long he has stayed at the previously visited artifacts. The same technology could very easily be used by an employer to trace an employee to see when he was at work and actually working, without the employee's consent. Privacy issues are therefore central to managing a user's context information, such as his location.

As described in Section 2, the Alipes platform offers applications a way to gain access not only to the local terminal's location, but also to other users' locations. It makes good sense to include support for privacy management in the platform, as trust and privacy issues may be central to the success and acceptance of location-based applications.

This section describes how the Alipes platform enforces privacy. However, transport level security and privacy are considered to be beyond the scope of this paper. The work of Alberto Escudero Pascal is a good source for more information on that topic [10], as well as the numerous IETF groups working on securing IP versions 4 and 6. Without transport level security it would always be possible to track a user or client using the MAC address, IP address or hostname of the terminal, regardless of any precautions made at the application level.

Note that secure links are required for this scheme to work. It would be quite simple to intercept messages and act as a middleman or simply fake being someone else. Some form of encryption and authentication is therefore required to ensure that the contracts and queries are transported across the network in a secure manner (especially as they are expressed in an easily readable text format using XML). Key distribution, authentication and encryption schemes like PGP [33] are therefore necessary, but will not be described in further detail here and are henceforth considered to be beyond the scope of this paper. For now we assume that all traffic is secure from attacks.

The focus of this work is also to define a privacy framework from the application programming view, thus not primary on how to graphically present and manage the framework in an application. However, presentation is very important in this context, but needs more investigation, for example through user studies, before any conclusions can be reported. Usability is therefore only commented on and not fully discussed in detail here.

3.1. Sensitive Information

Information about a user must be considered sensitive if it can be used to invade his personal integrity. An example of such information is the position of the user. If a user is only queried once, then the risk of violating his privacy is moderate, but if the user is continuously positioned and that data is logged and processed, then the data is highly personal. It can be used for criminal ends, such as to determine when the user is least likely to be at home in order to minimize the likelihood of a burglary being interrupted.

Other studies have shown that most users do not want complete automaticity of private data exchange, but instead want to be able to grant or deny any transfer of private data [22]. The user should therefore be able to fully control and limit any exchange of information. For instance, he should be able to limit the accuracy of a query (i.e. if the service does not require a very accurate answer then the answer should be adequate but not overly accurate) or to limit the number of queries.

He should also be able to turn positioning off completely in a simple manner, or to disable it at certain locations or in certain time periods or contexts. This means that a user should be able to limit positioning, for instance to working hours, to when he is located at work and when he is using a certain service. Note that turning positioning off completely means that all queries will automatically be rejected or denied, as the user at these times may feel that a query is intrusive and that may affect the user's trust in the system.

3.2. Design Criteria

The primary design criteria have generally been to keep the platform as simple as possible while creating a framework with sufficient support to enforce privacy. This means that:

- The framework should comply with national and international initiatives concerning legislation.
    - The user must at any time be able to turn off positioning completely.
    - The user must at all times be informed about positioning activities.

    - The user must be able to decide actively whether to grant or deny access from any party requesting location information.
    - The positioning activities must be logged for future reference.
- The set of rules and how they work should be easy to understand and process.
- The protocols and rules should be described in an easily readable format for debugging reasons.
- It should be easy to create user interfaces to manage the rules, while supporting different levels of user expertise or clusters of users.
- The system should also be functioning when a terminal is turned off, disconnected from the network, behind a firewall or using network address translation.
- The system should be designed to be distributed with as few central servers as possible and to allow peer-to-peer communication in real-time.

3.3. Limiting Access to Location Information

There are a number of different limitations on when and how other parties may locate a client, as expressed in the previous sections. Below are examples of limitations on when and how a certain party is allowed to position a client:

- Geographical area, e.g. only when in Luleå
- Accuracy, e.g. only that the user is within 1 km^2
- Time period, e.g. only working hours
- Number of queries, e.g. 20 queries per day
- User context, e.g. when using WaveLAN
- Usage, e.g. only for setting up conference calls
- Mutuality, e.g. 'you may position me if I may position you'
- Type of information, which is primarily the current location

Some of these limitations are nearly impossible to control. Even if a party is granted access only for a certain usage, i.e. to set up conference calls, there is nothing to stop that party also using the information for other purposes. Mutuality is another example of this, as one party could easily fake a location while another party reports its location truthfully. In the end these cases falls back on trust and on future legislative initiatives.

The Alipes platform defines a set of rules for limiting access to positioning information. These rules consists of an on/off rule, ban rules, generic criteria and contracts that are described in Sections 3.4 to 3.7.

3.4. On/off Rule

As previously stated in Sections 1 and 3.1, the user should be able to deny all access to his position in a simple way, independent of any general criteria or issued contracts (see Sections 3.5 and 3.6). The on/off rule will thus allow the user to easily turn off positioning completely, regardless of any other configuration.

3.5. Ban Rules

As in any communication scenario, there are times when you want it to be possible to ban all access by certain parties. The Alipes platform implements ban rules that are handled somewhat differently from other rules. They collectively act as a general criterion and are managed separately. The ban rules consist of a list of party IDs or IP addresses with connected time limits, which may be infinite. If a party is registered in the ban rule then a query or contract proposal from him is automatically denied. Note that general criteria and contracts will never be processed due to queries from a banned party; a banned party will simply be ignored.

A party can be automatically registered in the ban list if he repeatedly violates the system, but most commonly bans are registered by a user manually in the ban list or directly from the list of contracts. A user can also register bans using wildcards, such as *'*@*.luth.se'*.

If the ban rule is managed by a location server, then violators who are automatically registered can optionally be managed in a global rule for all users. This option can also be used for parties who are banned by a certain amount of separate users, or who are reported by single users to the location server administrator.

3.6. General Criteria

The general criteria define default limitations and apply to all parties. They can be of two different types: grant or deny. Granting means that general access is allowed under certain circumstances, while denying means the opposite. The limitations that are possible to specify globally are listed in Section 3.3. Note that deny criteria always takes precedence over grant criteria, unless they are nested.

3.7. Contracts

Agreements between parties concerning how location queries may be conducted are defined in the terms of contracts. Contracts describe the rights and responsibilities between two parties, a party being either a user or a location server (see Section 3.10 for more details on the location server). A party is identified by a unique user ID (basically any text string such as *user@host* if a user, or the service name or IP address in the case of a location server).

In other words, a contract defines when and how a certain party is allowed to position the client (see Section 3.3 for examples). A contract may contain grant and deny rules, and the deny rules have precedence over the grant rules unless they are nested. Each contract is also given an expiry date, on and after which they are no longer valid. Contracts are thus identical to a set of general criteria, with the distinction that they are only valid for a certain party.

3.8. Query Strategy

Figure 5 depicts the query strategy applied when a client is queried for information. The client will first check the on/off rule and thereafter the ban rules to see if a deny reply could be sent directly.

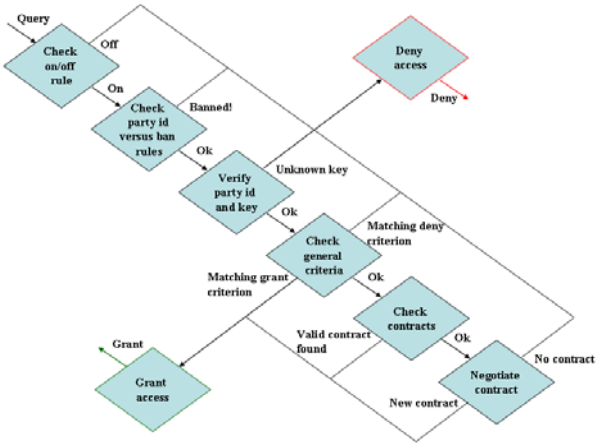

Figure 5. Query strategy.

If the on/off rule is ‘on’, then the client will verify the party identifier and the key in order to authenticate the querying party. If no valid ID/key pair is found then a deny reply will be sent.

The next step is to find matching deny criteria, which if found will result in a deny reply from the client. If no matching deny criteria are found but there are matching grant criteria, then the client will grant the party access to the location information.

If no matching general criteria can be found, then the client will seek a valid contract. If a valid contract exists, then the client will again grant access and reply by sending positioning data back to the querying party.

If no valid contract can be found, then the client will reply with a prompt to set up a new contract, at which point a contract negotiation will be initiated. If a new contract can be negotiated, then the client will grant access as if a valid contract was found, and if not then a deny reply will be sent.

The querying party can also directly propose a new contract and thus initiate a negotiation before querying for data (if, for instance, there has been no previous communication from that party to that particular client). No search for a valid contract will then take place, but the rest of the strategy as described above will be carried out.

Note that the Alipes platform, when granting access to location information, allows for ‘keep-alive’ sessions as well as single replies. The keep-alive session always check if the session is still valid before sending any data by going through the steps described above, with the exception of authentication and contract negotiation. If the keep-alive session is terminated, then the party needs to initiate a new query including contract negotiation.

3.9. Local Rule Management

The user can select whether or not he wants to be notified when a contract proposal is received and no general criteria match or no valid contract can be found. He will then be prompted to accept, modify or deny the contract, as well as have an option to add, remove or modify general criteria or ban rules. If a contract proposal is received and there are matching general criteria, then the client will deny the contract while notifying the party that queries can be sent that match the general criteria.

All the criteria and contracts are maintained in a list, from which the user can select to view statistics of their usage. He can see when, how many times and by whom he has been positioned. The contracts and criteria can be enabled and disabled one by one. There is also one general on/off rule, which can be used to disable positioning completely.

3.10. Location Servers and Remote Rule Management

A location server is an intermediate party, as stated in Section 2.3, which can be granted rights to act on behalf of a user. A user can therefore delegate the right to manage his position information under certain circumstances. These rights are handled by setting up rules at the location server, which then acts as a negotiator for contracts for other parties. Any query or contract proposal sent to the client will then be redirected to the location server.

Using a location server will automatically disable management of local general criteria and contracts in the client itself, as the client will delegate these rules to the location server. The only visible difference to the user is that he can select local or remote management of rules (which for some clients may be pre-selected and not visible). The rules at the location server can, however, be managed at any time by the user, in a way identical to the local management of rules in the client. A user can select whether or not he is to be notified by the location server when no matching general criteria or no valid contracts can be found, just as when managing rules locally.

3.11. Meeting Design Criteria

The intention was to keep the design of the platform as simple as possible while enforcing privacy sufficiently. The set of rules have thus been designed to be simple to understand and process. An alternative design would be to not separate the different rules and allow for any mix. This design chosen should make the process easier to understand, as it is split into distinctive parts. The query strategy should therefore also be simple to process and understand.

Rules and protocols are expressed in XML and each query is logged to enable easy debugging and monitoring during development.

An application can use the Alipes platform and the privacy interfaces in a number of ways to best meet the users' and the applications' needs and requirements. For instance, it is simple to create user interfaces supporting different levels of user expertise. The application can offer minimal management of the rules by handling rules automatically or deferring decisions by default to a location server, so that a novice user basically can merely turn positioning on or off. The application can also offer limited management of the basic rules for intermediate users, where for instance nested rules are excluded for simplicity. The application can furthermore allow complete control for advanced users.

Note also that the previously identified clusters of users may take advantage of the privacy mechanisms detailed in this paper. As stated by Ackerman et al, the cluster of privacy fundamentalists and the cluster of marginally concerned users may find extremely simplified interfaces to be adequate for their purposes. A marginally concerned user would only need to specify those few already constrained instances in which he would not permit information collection practices [22]. These two clusters could thus make use of a few basic not-nested rules or prepared composite rules. However, Ackerman et al also state that the pragmatic majority of users will require more sophisticated and varied interface mechanisms to be most at ease. These pragmatic users would typically employ many strategies across a wide range of finely weighed situations and would thus be likely to require full access to the privacy mechanisms, such as nested rules, as it is unlikely that a highly simplified interface will satisfy them.

The involvement of location servers allows the system to be used when a terminal is not directly accessible, for example when behind a firewall or using network address translation.

Due to the limited availability of IP-multicast and mobile IP, the system could initially not be designed to be fully distributed. Instead a topology of servers is used in order to look up a terminal's IP address. It is however possible to communicate peer-to-peer in real-time.

3.12. Implementation Details

The first implemented prototype application using the location server and context information is the Tracker application depicted in Figure 6. The users Joe and Katja are busy at the moment, which can be seen from their red position markers. PB on the other hand is available and this can be seen from his green position marker. The user's current location (Kåre) is marked with a yellow marker.

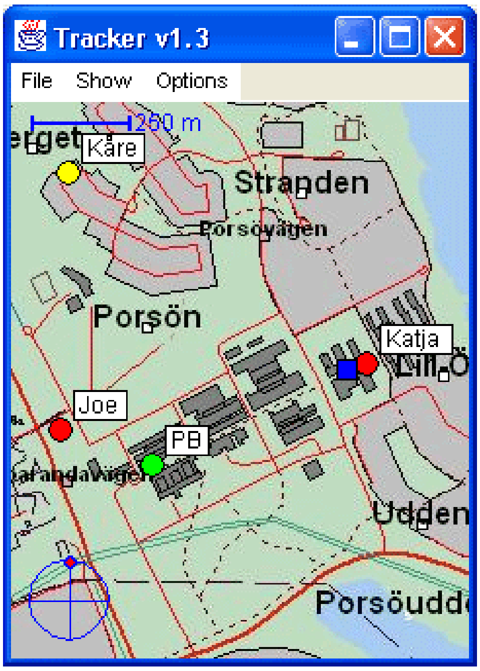

Figure 6. The Tracker Application.

The privacy parts of the Alipes platform are currently being implemented, optimized and tested. The platform is implemented with Personal Java 1.1.3 and is intended to run on a StrongARM PocketPC with the SUN Personal Java 1.1 beta or the Insignia Jeode Personal Java 1.2 virtual machines.

We have chosen to base our first implementation of the user location look-up mechanism solely on a scheme similar to SIP, because it is a simple and proven solution that does not have the deployment problems associated with mobile IP or IP multicast (see Section 2.2).

The location queries and all the exchange and storage of contracts and criteria described in this paper are implemented using XML.

All the messages are encrypted using a simple secret key implementation, where the secret keys are distributed together with the party identifier (i.e. user or server ID, see Sections 2.2 and 3.7). Each party identifier is thus paired with a secret key to allow authentication. Note that the exchange of party identifiers and secret keys needs to be done in a secure way.

3.13. Related Work

A great deal of work has already been performed regarding privacy issues in computer systems. Many of the existing standardization groups, like LIF, WAP Forum or the Location Positioning Workgroup [11-12], are strictly restricted to members only (for which reason comparison is incomplete).

The IETF Geographical Location/Privacy Workgroup [13-14] is very interesting and directly related. The World Wide Web Consortium (W3C) and their Platform for Privacy Preferences Project (P3P) enable Web sites to express their privacy practices [34]. The specification is relevant to the present work, but does not directly apply since it is designed for the web and not for mobile applications in general. It does compare in part at a technical level, as similar technologies are used for implementation.

There is also related work within the field of human-computer interaction, such as by Westin and Ackerman et al [1,3,21-23], where the user studies are of special interest and can be used when designing the privacy mechanisms themselves or applications based on these privacy mechanisms using the Alipes platform.

4. Future Work

The next logical step to drive future innovation forward is to conduct user studies to establish the current practice and a use-driven development methodology. We will follow the research methodology exemplified in Carroll et al [35], which synthesizes theory from multiple data sources (including questionnaires, observation, diarizing and interviews). This will provide a basis for envisioning future applications by using contextual scenarios, acting out and participatory design workshops [36]. Of particular interest are the following questions:

- How do users feel about the fact that they may be positioned? (With different levels of access to and control of the privacy mechanisms.)
- Is the management of privacy rules simple and understandable enough for the average user?
- What level of management of the rules is necessary and used? (Are nested rules really necessary, as they add a great deal of complexity?)
- How many rules will a user define in general, and in contrast to how much the applications are used in particular?
- How do the users' rules generally interact?
- How can the management of position information be extended so that it will also include context information in general.

Furthermore, how do context-aware systems affect the non-users? Herstad et al suggests in [37] that also persons not using the technology must be considered. For instance, one person may be positioned when together with another person. What are the effects of also involving the non-users into the picture? Should persons related to each other jointly decide on common rules?

Finally a more sophisticated scheme for authentication, encryption and key distribution is necessary, as the really rudimentary secret key mechanism currently implemented has too many limitations and is not easy to use.

5. Summary and Conclusions

It can be concluded that any service that handles private information must be protected by privacy safeguards and that the users of the service must be fully informed of the purposes of any usage of their private information and must have the right to determine the use of that information. The users must also be able to deny or disable use of such information at any time. It is clear that the user is the owner of private information and therefore must be given the right to give informed consent of any use thereof.

The Alipes platform enables a wide range of new services that use private information, such as the context and location of the users. The platform therefore includes a possible solution for managing the distribution of positioning information. The platform defines an on/off rule, ban rules, general criteria and contracts for managing queries from external parties, such as users or location servers. It also uses a simple method to authorize parties using ID and key pairs.

A user can view the available rules and modify them as he deems fit. New proposed contracts may be granted, denied or modified, and there is also an option to also add, remove or modify general criteria. There is also a way to ban use by certain parties for a certain period or to disable positioning partially or completely. Applications based on the platform can be designed to support different levels of complexity when managing privacy rules, from a very basic set for new users to more complex nested rules for advanced users. It also supports users with various levels of consciousness about their privacy.

The platform also includes of location servers, which are delegated the right to manage a users personal information and thus act as intermediaries. This also allows a user to be positioned even if his terminal is turned off, not connected to the Internet, behind a firewall or using network address translation.

The novelty of this paper lies in how privacy is enforced through the Alipes platform and how the design supports distributed real-time applications in mobile terminals.

References

[1] Alan F. Westin, “Privacy and Freedom”, NY: Atheneum Press, New York, USA, 1967.

[2] M. Köhntopp, A. Pützmann, “Anonymity, Unobservability, and Pseudonymity - a Proposal for Terminology”, Information Hiding Workshop, Pittsburgh, USA, April 2001.

[3] M. Ackerman, T. Darrell, D. Weitzner, “Privacy In Context”, The Journal of Human-Computer Interaction, Special Issue on Context-Aware Computing, Volume 16, numbers 2-4, 2001.

[4] Lagtext, Personppgiftslagen, April 2002, http://www.notisum.se/rnp/sls/lag/19980204.HTM.

[5] Pul.nu, April 2002, http://www.pul.nu/.

[6] European Council and Parliament, The European Union Directive 95/46/EC concerning the protection of individuals with regard to the processing of personal data and on the free movement of such data, http://www.privacy.org/pi/intl_orgs/ec/eudp.html.

[7] European Council and Parliament, Data Protection Telecommunications Directive 97/66/EC of 15 December 1997 concerning the processing of personal data and the protection of privacy in the telecommunications sector, http://europa.eu.int/ISPO/infosoc/telecompolicy/en/9766en.pdf.

[8] European Commission, Proposal for a new directive concerning the processing of personal data and the protection of privacy in the telecommunications sector, http://europa.eu.int/comm/information_society/policy/framework/pdf/com2000385_en.pdf.

[9] European Parliament, European Telecommunication New Regulatory Framework, Proposal concerning the processing of personal data and the protection of privacy in the electronic communications sector, 12 July 2000, http://europa.eu.int/information_society/topics/telecoms/regulatory/new_rf/index_en.htm.

[10] A. Escudero, “Privacy in the Next Generation Internet”, Dissertation Proposal, Stockholm, December 2001, http://www.it.kth.se/~aep/publications/phd-proposal/.

[11] Location Inter-operability Forum (LIF), 2002-02-24. http://www.locationforum.org/.

[12] WAP Forum, April 2002, http://www.wapforum.org/.

[13] IETF Geographic Location/Privacy (geopriv), April 2002, http://www.ietf.org/html.charters/geopriv-charter.html.

[14] J. Morris et al, “Framework for Location Computation Scenarios”, IETF draft, March 2002, draft-morris-geopriv-scenarios-01.txt.

[15] James Nord, Kåre Synnes, Peter Parnes, "An Architecture for Location Aware Applications", HICSS-35, Big Island, Hawai´i, USA, January 2002.

[16] Electronic Privacy Information Center, 2002-02-24, http://www.epic.org/.

[17] Privacy International, 2002-02-24, http://www.privacyinternational.org/.

[18] Privacy.org, 2002-02-24, http://www.privacy.org/.

[19] Consumer Privacy Guide, 2002-02-24, http://www.consumerprivacyguide.org/.

[20] Center for Democracy and Technology, on Data Privacy, 2002-02-24, http://www.cdt.org/privacy/.

[21] A. Westin, “E-commerce & Privacy: What Net Users Want”, Hackensack, NJ: Privacy & American Business, 1998.

[22] M. Ackerman, L. Cranor, J. Reagle, “Privacy in E-Commerce: Examining User Scenarios and Privacy Preferences”, ACM Conference on Electronic Commerce (EC'99), Denver, Colorado, USA, November 1999.

[23] A. Westin, Harris-Equifax Consumer Privacy Survey, GA: Equifax Inc, Atlanta, USA, 1991.

[24] Peter Parnes, Kåre Synnes, Dick Schefström, "mStar: Enabling Collaborative Applications on the Internet", Journal of Internet Computing, September/October 2000

[25] S. Björk, J. Falk, R. Hansson, P. Ljungstrand, P. “Pirates! - Using the Physical World as a Game Board”, Interact 2001, IFIP TC.13 Conference on Human-Computer Interaction, July 9-13, Tokyo, Japan.

[26] Ericsson Mobile Positioning System (MPS), April 2001.

[27] M. Nilsson, J. Hallberg, “Positioning with Bluetooth, IR and RFID”, Master’s Thesis, Luleå University of Technology, February 2002.

[28] Open GIS Consortium, March2002, http://www.opengis.org/.

[29] Charles Perkins, Mobile IP Tutorial, March 2002, http://www.computer.org/internet/v2n1/perkins.htm.

[30] S. E. Deering, “Multicast Routing in a Datagram Internetwork”, Ph.D. Thesis, Stanford University, December 1991.

[31] M. Handley et al, “SIP: Session Initiation Protocol”, March 1999. IETF RFC-2543.

[32] W. Yeong, T. Howes, S. Kille, “Lightweight Directory Access Protocol”, March 1995. IETF RFC-1777.

[33] The International PGP Home Page, March 2002, http://www.pgpi.org/

[34] The Platform for Privacy Preferences 1.0 (P3P1.0) Specification, W3C Proposed Recommendation, 28 January 2002, http://www.w3.org/TR/P3P/.

[35] J. Carroll, S. Howard, F. Vetere, J. Peck, J. Murphy, “Just what do the youth of today want? Technology appropriation by young people.”, In R.H. Sprague (ed.) Proceedings of the 35th Hawaii International Conference on System Sciences (HICSS-35), Hawai’i, USA, January 2002.

[36] S. Howard, J. Carroll, J. Murphy, J. Peck, F. Vetere, "Provoking Innovation: Acting-out in Contextual Scenarios", BCSHCI, UK, September 2002.

[37] J. Herstad, D. Stuedahl, D. van Thanh, “Non-user Centered Design of Personal Mobile Technologies”, In proceedings of IRIS 23, Uddevalla, Sweden, August 2000.

Acknowledgments

Thanks are extended to Arkady Zaslavsky at Monash University, Jennie Carroll and Steve Howard at Melbourne University, as well as Magnus Erixon and Anette Järn at Telia Research for their comments, suggestions, and feedback.