

Localization to Enhance Security and Services in Wi-Fi Networks under Privacy Constraints

Gareth Ayres¹, Rashid Mehmood¹, Keith Mitchell², and Nicholas J.P. Race²

¹ Civil and Computational Engineering Centre, Swasea University, Swansea SA2 8PP, UK

² Computing Department, InfoLab21, Lancaster University, Lancaster, LA1 4WA
{g.j.ayres,r.mehmood}@swansea.ac.uk,
{k.mitchell,n.race}@lancaster.ac.uk

Abstract. Developments of seamless mobile services are faced with two broad challenges, systems security and user privacy - access to wireless systems is highly insecure due to the lack of physical boundaries and, secondly, location based services (LBS) could be used to extract highly sensitive user information. In this paper, we describe our work on developing systems which exploit location information to enhance security and services under privacy constraints. We describe two complimentary methods which we have developed to track node location information within production University Campus Networks comprising of large numbers of users. The location data is used to enhance security and services. Specifically, we describe a method for creating geographic firewalls which allows us to restrict and enhance services to individual users within a specific containment area regardless of physical association. We also report our work on LBS development to provide visualization of spatio-temporal node distribution under privacy considerations.

Keywords: Location-Awareness, Security, Privacy, Visualisation, Wireless Networks.

1 Introduction

Mobile services are increasingly pervasive due to the recent developments in localization, context-aware, and mobile technologies. In today's modern society, users are expecting location based and context-aware services to aid them in their daily lives. The increased popularity and usage of location-aware and pervasive mobile systems have brought two broad challenges. Firstly, access to wireless systems is highly insecure due to the lack of physical boundaries. Secondly, the inherent mobile nature of users within wireless networks provides an opportunity to offer location or context based services, however, it does render the users vulnerable since highly sensitive information about them can be extracted from location based systems. First problem can be addressed by developing security systems which exploit node location and contextual information. Security could be considered a service too and hence systems to provide both security and services can be designed and developed on a generic level such that system security is ingrained into location based services and is used as an inherent tool to protect user privacy.

In this paper, we present our ongoing work on developing systems which exploit node location information to enhance security and services under privacy constraints. The work is based on the initial results of two complimentary methods developed to track node location within large production (or live) University networks comprising of large numbers of both fixed and mobile clients. Lancaster University has mainly focussed on (finer-grained) indoor tracking while Swansea's focus has been on (course-grained, building level) outdoor localisation. An aim is to investigate whether useful contextual information can be extracted from a production wireless network in order to offer value added services.

Specifically, we describe a method for creating geographic firewalls based on user location enabling the restriction and/or enhancement of services to users within a pre-defined 'containment area'. Crucially, we are able to prevent access to mobile users whom are connected through the same Access Point as other users in neighbouring rooms without degrading the service provided to legitimate users. The motivation for this stems from a number of institutional requirements such as being able to disable wireless access to specific rooms or lecture halls on campus during exam periods without affecting neighbouring rooms. Furthermore, we present our work on developing location based services using visualisation technologies for campus network at node activity level. Such location based services (LBS) could be used, for example, by the network operators to find and predict usage patterns, congestion points, traffic and network growth. This strand of work uses open source solutions including Google API to provide visualization of spatio-temporal distribution of network nodes. We describe our work to date and the careful considerations made to user privacy, and show how we are able to extract useful information about nodes activities including mobility across campus without affecting user privacy.

In the process of describing location based firewall, visualisation and other LBS which we mentioned in the paragraph above, we give a discussion on the design and architectural choices for systems that provide location based security and services. The rest of the paper is organised as follows. Section 2 gives a brief literature review of location based security and services. Section 3 describes the network architecture for Swansea and Lancaster University. Section 4 describes the design of location based systems providing security and services. Section 5 presents and discusses results: results on location based security and visualisation services are discussed in Section 5.1, while Section 5.2 gives a brief literature review and discussion on Privacy in the context of Swansea's work on network visualisation. Section 6 concludes the paper. This work is part of the Janet UK Location Awareness Trials [1] which aims to explore the possible applications of location awareness in a wireless context within the education sector.

2 Location Based Security and Services

Wireless and mobile technologies have progressed to a point where the use of wireless networks is common place in most businesses, academic institutions and homes in modern societies. The move from wired to wireless communication systems has brought about many challenges as well as opportunities. Location Based Services (LBS) are one such opportunity. LBS apply to wide range of applications that exploit

the physical location of the user through GPS or other technologies to facilitate user-specific and personalized services. On the other hand, users are no longer bound to physical locations and are free to roam wherever there is sufficient wireless coverage. This ability of users to roam and become 'mobile' has brought about lots of new hurdles and problems with respect to security and access control. It is no longer sufficient to rely on the assumption that users who connect to the network are users who are already physically located inside the institution. Wireless networks penetrate physical boundaries and have therefore required a new approach to security.

The problem of securing the traffic over the wireless medium has been addressed with varying success over the years with encryption protocols such as WEP, WPA and WPA2. These are now addressed by the IEEE 802.11i standard. IEEE 802.1x allows for enterprise level authentication over wireless networks and its use combined with IEEE 802.11i currently provides a secure and accountable solution to wireless network security. This solution, however does not address the problem of providing different levels of network access control and security according to the location of the user. To achieve location based security it is important to be able to accurately obtain the location of the user. The problem of localization of wireless nodes has been tackled using a number of wireless technologies. A fundamental paper in this field is [2] which identifies a number of techniques for locating and tracking wireless nodes using RF signal measurements. Outdoor location discovery is usually solved with GPS, but this approach will not work indoors. A number of technologies have been explored to help with indoor location discovery. The Active Bat [3] system uses ultrasonic RF measurements to locate nodes, while the Active Badges [4] system uses infra-red. The Cricket [5] system uses a combination of ultrasound and RF beacons. These technologies are surveyed and compared well in [6, 7]. See also our earlier work for a brief review of LBS applications and methods, and UWB (Ultra WideBand) based localisation in indoor environments [8].

Although the technologies mentioned work well they are aimed at specific problems and require significant hardware and/or software investment. The more specific problem of the localization of IEEE 802.11 nodes has also been addressed by a number of researchers. Nearly all localization techniques are based around the RF signal measurement techniques originated from [2] and make use of the received signal strength indicator RSSI value. The initial problems identified with Wi-Fi network localization are that it is not very accurate (up to 1-2 meters accuracy at best) and requires an understanding of the physical layout of the wireless environment. The problem of improving accuracy from RSSI values has been worked on by [9-12] and a number of algorithms have been proposed with varying results. The problem of the wireless physical environment is one that has been 'passed-off' by a number of commercial applications which put the burden of calibration on network administrators. This approach is time consuming and tedious, and is often not desirable to network administrators and as a results does not get done correctly or even at all. A number of solution to this have been proposed, one of the first being [13] which work towards a easy self calibration of a wireless environment with limited user input. There are localization techniques not based on signal strength values but on data obtained through RADIUS packets and SNMP data. These were first described by [14]. There are also a number of attack vectors possible based on location information, these are described well in [15].

Once the location of a node has been determined, it is then possible to build location based security into the system. There are a number of ways to achieve this which are either client (node/user) side, server (network infrastructure) side or a combination of both. A client side example based on user intervention is [16], it relies on users selecting their location via a GUI once they have associated to a wireless network. This work was the precursor for the location selection option seen today in Microsoft Vista once user first connects to a wireless network. A server side example is [17] which has a server that monitors RSSI values of clients and then modifies firewall rules accordingly. The approach adopted by a lot of commercial products generally involves a similar server side approach, making use of RSSI values and similar location discovery techniques as described in [2]. Most products then augment the location data with user inputted maps and floor plans to provide LBS such as asset tracking and usage visualization. Privacy is an increasingly important and sensitive topic today, which we discuss in Section 5.2, in context of location based services.

3 Wireless Infrastructure

The Lancaster University campus is set in 250 acres of parkland and lies approximately 3.5 miles south of the City of Lancaster. The University operates a wireless network both on campus and within areas of the city of Lancaster. Overall, approximately 400 Access Points have been deployed with the aim of providing coverage across all academic departments and service buildings (i.e. libraries) as well as social and public spaces. Wireless network coverage is not provided to student residences since access within student rooms is provided through a separate 100Mbps wired network infrastructure known as RESNET. The hardware infrastructure deployed at Lancaster is based on Cisco's Wireless Location Appliance (LA). The overall solution comprises of the Cisco LA which is used in conjunction with the Cisco Wireless Control System (WCS), Cisco wireless LAN controllers, and Cisco Aironet lightweight access points (LWAP). The APs are essentially used to record sightings of wireless clients which is stored in the LA database. The WCS provides a web front end for information visualization and management of the overall system. Additionally, the location information is also available to third party applications through a Web Services (SOAP/XML) interface (API) on the appliance and it is through this interface which we extract mobile node information.

Swansea University is largely a single campus oriented university based on the Swansea sea front consisting of 43 Buildings surrounded by Singleton Park. There are also two areas off campus used for halls of residence. Swansea University provides an IEEE 802.11g wireless internet service to students and staff within all of its campus buildings, halls of residence on campus, Student Village off campus and Beck House residence. The wireless network is primarily composed of 760 Cisco 1100/1200 Aironet Access Points in Light Weight mode, 4 Cisco Wireless Service Module (WiSM) blades and a Cisco Wireless Control Server (WCS).

4 Location Based Security and Services – Design and Architecture

As mentioned earlier, wireless networks cannot be secured using physical bounds alone since any device within range of a wireless signal is able to listen or attempt to connect to the network. Also, in certain environments, such as a university lecture

theatre, it may be necessary to restrict certain network traffic types to users confined by the physical bounds of the lecture theatre, while still allowing users outside the theatre to connect through the same wireless network access point. Location based security approaches can help solve these problems. Security in this context can be considered a service, leading to the development of generic system architectures to provide both security and services based on location or context. This section presents and discusses the design and architecture of our systems providing locations based security and services. Section 4.1 discusses the system design in the context of location based firewall (security) being developed at Lancaster. Section 4.2 briefly presents the work on location based services and visualisation at Swansea University. From a system design perspective, most of the architectural details of the two strands of work at Swansea and Lancaster are similar.

4.1 Toward a Location Based Firewall at Lancaster

There are two distinct challenges to providing location based security (or service). The first is how you ensure you receive timely updates regarding user location in order to rapidly resume or restrict network access (or to provide another service). The second relates to the method of controlling (i.e. restricting/resuming) network access in real time for production environments in situations where mobile devices are highly mobile. Therefore, in order to effectively deploy a location based firewall solution within an existing wireless infrastructure, the development of additional services are required, namely: *Location Data Gathering* to collect and manage sighting pertaining to digital assets (tags, mobile devices, etc); *Location Based Security Policies* in order to accurately link geographic location with authentication, authorization and accounting; *Access Control* in order to implement and enforce the chosen security policies. These are discussed in the following three subsections.

4.1.1 Location Data Gathering

In this section, we introduce the method adopted for obtaining and determining user location within our production wireless network. We (at Lancaster) have developed third party software which communicates with the Cisco Location Appliance (LA) through a Web Services (SOAP/XML) interface (API) on the appliance. This interface exposes details relating to all aspects of the wireless infrastructure (maps, clients, statuses, etc) and these sources of information are used to extract mobile node information in a timely fashion. There are two basic approaches to collecting this information. Polling: Here, a third party application polls the location engine via SOAP/XML and gathers a list of clients and their current location. Notifications: Using this method, the location engine is pre-configured in order to generate SNMP traps or push XML to a configured IP address when pre-determined events occur, such as a client moving into a defined region.

The software developed at Lancaster for the Location Awareness trials has been written in C#.NET and makes use of the Cisco Location Service API (via XML Web Services) in order to enable third party applications to access the location information stored within the location appliance (LA). Our initial implementation consists of two components, the first manages the communication to/from the location server and gathers the relevant location information and the second is responsible for managing the access control. For the purposes of application described below, we poll the

location appliance in response to specific user requests rather than having a series of pre-defined notifications. Once data of interest has been retrieved from the LA, this is processed by the access control component which makes explicit requests to the location based firewall, which in turn controls access to the wireless medium.

The data gathering engine consists of a dedicated process which communicates with the Cisco Location Appliance server using SOAP/XML over HTTP or HTTPS. This communication protocol consists of 3 basic message types: **Request**: Sent from a client to the server to set or get information stored in the LA. **Response**: Sent from the location server to a client in response to a particular request. **Notification**: Sent from the server to a client asynchronously upon a particular server event occurring. A client application may first register for notifications based on a set of criteria the client application defined in the request for notification. With the exception of the initial login request all subsequent SOAP request messages consist of a *request name* to identify the method being called on the server, a *unique session id* (the server rejects all calls not containing a valid session id), a list of *objects* that act as the parameters to the method request and an optional list of *attachments* (for example images). All server response messages are indicated by a response flag, followed by a list of objects for the client to process, which contain the results of the request.

4.1.2 Implementing Location Based Security

The application we have developed as part of this paper has been designed to allow lecturers or speakers to restrict wireless access within rooms for a specific time period. This stems from a series of institutional requirements such as being able to disable wireless access to specific rooms or lecture halls on campus during exam periods without affecting neighbouring rooms. In practice, this can also be applied to restricting access during lectures and conferences on campus. In essence, our aim is to link geographical location with authentication, authorization and accounting (AAA). At present corporations like Cisco do not link these two elements and it is here where some of our efforts have been focused in order to bridge between them.

The prototype client application we have developed resides on each of the lecture room PC's (each lecture hall on campus has dedicated data projector and networked PC). An authenticated user logged into the PC (i.e. members of staff only) is able to navigate to a simple web page which, based on their location, provides a list of rooms available in that building and on that floor. From this page the user is able to select a room and a time window before selecting the disable button. Initiating the request to disable access initiates a call to the third party location and tracking service (LocoTrak). The LocoTrak service acts as a proxy service and handles all requests targeted for the Location Appliance (LA). Upon receiving a request, LocoTrak queries the LA for a list of all mobile devices stored within the database located at that location at that time. The LA returns a list of *AesMobileStationLocation* objects filtered by the specified room id. The result is a list of MAC addresses recorded for that location and, by default, the location appliance returns the list of devices sighted within the past 24 hours. This list is then filtered by LocoTrak so that only those recorded in the past 2 minutes are stored. This MAC address list is then forwarded to the GeoFirewall server, which is responsible for granting/denying access to the network.

The LocoTrak service provides a response to the user's web browser and indicates the number of initial devices found. Further, the service spawns a new process which

runs for the duration of the specified time period, for example, a 50 min lecture. During this time, the thread periodically (5 min intervals) queries the LA the list of all mobile devices stored within the database located at that location. New devices discovered are forwarded to the GeoFirewall to deny access temporarily; similarly, devices who appear to have left the area (after not been discovered after 3 continuous scans) have their access restrictions lifted.

4.1.3 Access Control

The process of access control is non-trivial and within a production environment we need to have a reliable way of receiving and processing location updates. The first step is to determine whether any restrictions are necessary based on the individual users' current locality or whether they are allowed access to the network. We believe that there are four different approaches to dynamically restricting access to network services and we are currently investigating each in order to evaluate the most effective and reliable solution.

a) Preventing clients from connecting to the wireless network by setting up MAC filters on the Wireless LAN Controller and by sending de-authentication requests to clients: From a conceptual point of view this appears to be the most trivial method and would appear to offer the most efficient and fast method overall because mobile clients would become de-authenticated and lose network connectivity immediately. However, the main challenge is that on its own it is not effective enough as it could be bypassed by changing the MAC address on the mobile client device. A further drawback of significant concern within a live production environment is that this method would provide a distinct lack of feedback to the end user which would detail the reason why they have lost network connectivity. For this reason, we have not pursued this avenue further.

b) Use our existing role based firewall and captive portal hardware from BlueSocket: We use a BlueSocket BSC-2100 to restrict access to network services and to limit bandwidth usage. This is done based on roles that the BlueSocket assigns to users during the authorization phase. Roles are based on faculty membership and on user status, such as undergraduate, postgraduate or staff. The BlueSocket also supports the ability to quarantine individual users based on their MAC address and, further, provides a SOAP interface which makes it flexible enough to allow third party applications to access its features. However, during the course of our initial trials we have discovered that although the quarantine function when applied does accurately block clients immediately, the re-establishment process is extremely cumbersome. More specifically, once a mobile client (e.g Windows XP, Vista, Linux) has their quarantine restriction lifted, they do not automatically gain access to the wireless medium as one might expect. To re-gain access, a client has to re-authenticate with the wireless network explicitly and to best achieve this, the wireless interface ideally needs to be disabled and then re-enabled in order to initiate the re-authentication process. As with the previous option, this means that behaviour is likely to be inconsistent and unexplained from the end user point of view making this approach unsuitable for production services. Additionally, since the quarantine is again based on MAC address only, one could easily bypass it.

c) *Dynamically configure a traditional firewall*: An alternative approach we are investigating is the use of IPTables on a dedicated Linux based server in order to restrict access. This approach can be based on client IP address or physical MAC address, which although could both be manually changed by users and therefore bypassed, has other advantages. Other issues which we seek to explore fully are, whether this method is able to block established streams effectively (e.g. a client already streaming a video prior to being blocked) and, how well this approach would scale given a situation where hundreds or thousands of clients are connected. The biggest advantage of this approach is that the inclusion of another entity means that we can redirect web (HTTP) traffic to a well known web page which can be used to provide clear feedback as to why a user has been blocked as well as support instructions.

d) *Null route clients at Layer 3*: A final involves null routing each discovered IP address at the border of the network and we do not currently have the required hardware in order to test this approach fully.

4.2 Location Based Services

We now briefly describe the work being carried out at Swansea on developing Location Based Services using open source software and technologies; referred to here as the Locaware System. The Locaware system is composed of two servers and applications written in Java and PHP. The first server 'locaware' acts as a listener to collect all the location data sent to it by the Cisco WiSM's. Once it receives this data - in the form of SNMP Traps - it records the data to a database. PHP scripts running on the web server can then be called via a 'HTTP GET' API call to analyse the data in the database and provide location information. This can be seen in Figure 1. PHP was used to define groups of access points, and then perform localisations calculations based on the data contained in the location database. In the example in figure 4 each group represents a building. Node1 makes a request for a LBS embedded in a web page. The web service will use an API provided by the Locaware system to request the location of Node 1 using the IP address from the HTTP headers. This IP address is then looked up in the DHCP server for the corresponding MAC address, and hash code is generated. This hash code is then used to look up current location of the node from the location database.

This system will allow for the easy adoption of the use of location based services by Swansea University web service developers. The number of possible location based services feasible is considerable, but current developed are: Location aware wireless statistics and reporting allowing users to make informed decisions on where is the least congested place to access the wireless network and enhanced wireless problem reporting that allows technicians to provide better support. In particular, we are developing network visualisation technologies, as a location based service, to aid us in traffic and network monitoring, prediction of congestion points, future traffic and network growth. The need for such tools arises because wireless network deployments are now a fundamental service provided by universities and large institutions alike, and a more detailed understanding of the usage patterns of users is

vital in order to provide an acceptable level of service. With the growing trend of wireless deployments performing a uniform distribution, or ‘blanket installation’, of wireless coverage across a campus, emerges the inherent problem of congestion. As users do not spread themselves geographically uniformly, but tend to group together in locations such as eating areas and lecture rooms, a blanket installation may require the addition of extra wireless coverage in areas of high congestion. Some Preliminary visualisation results are discussed in the next section.

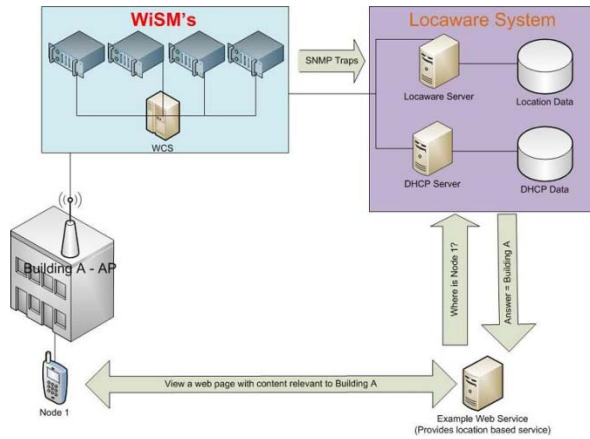


Fig. 1. Location Data Gathering

5 Results and Discussion

5.1 Network and Node Activity Visualisation

We now present and discuss some selected results, first for Lancaster and then for Swansea. Figure 2 shows the mobile node activity across the campus for a 24-hour period on 30th March 2009, which reflects a typical day. We have not yet carried out individual measurements per location on campus but know that the most heavily used areas on campus are the Library and the Management School, since these have the greatest density of APs, meeting rooms and work spaces available. Figure 3 shows a heat map for one of the floors within a building on campus, InfoLab21. This map shows the location of the APs and the coverage area supported by that AP, all associated clients and their IP addresses and finally the regions (or containers) defined in relation to that floor.

The key results from Lancaster’s perspective relate to the quality and accuracy obtained from the Cisco Location Appliance, specifically: **Incompleteness**: The current version of the Cisco LA SOAP API does not provide data consistent with similar queries entered via the WCS. Data appears to be inconsistent internally when querying for the same object using both the web interface and the XML based API. **Scale**: Location values (x, y coordinates) mismatch. **Determining and Managing location**: The Cisco Location Appliance does accurately present the third party developers with location information which can be related to the granularities of buildings, rooms and

containment areas. Associated with each response for a location is a confidence factor which can be used to determine the reliability of the response. A large confidence (say 75) factor implies that the system has calculated with 95% accuracy that a node is located within a 75m² region. Further, these results can be filtered according to a number of factors such as SSID associated and specified time periods.




Fig. 2. Node Activity across campus




Fig. 3. Heat map for a floor (InfoLab21)



Fig. 4. Wireless Usage at 14:00




Fig. 5. Wireless Node Activity

In general, we found evidence to support the general guidelines (provided by Cisco) that each mobile device or asset needs to be heard at better than -75 dBm by at least three access points or monitoring stations in order to provide a useful and accurate location. Although it is clear that the more APs/Monitors that hear a device the better, this obviously depends on the deployment environment and building layout(s). Optimal AP/Monitor density is approximately one every 25 meters depending on the environment and WLAN requirements and in general, placing APs towards the perimeter of rooms and coverage areas provides the best overall coverage and location fidelity. Additionally, for providing coverage within corridors or walkways, the best results are obtained when the APs are staggered along alternating walls rather than placing the APs on the ceiling along the centre line of the corridor. Despite this, we have also found that a large number of deployed access points do not guarantee the correct density needed for effective location awareness. Specifically, the placement of access points for the provision of a wireless network service doesn't mean that this automatically achieves good location fidelity and location awareness possibilities.

We now discuss visualisation of wireless nodes activity for Swansea. Figure 4 shows the wireless usage for a day and time chosen randomly (Monday 2 June 2008 at 2pm). The figure shows a total of 480 node associations with users spread around most of campus, and still heavily at the halls of residence. The Google API was used to display a Google Maps visualization of Swansea University campus using satellite imagery. The usage data was then used to create an overlay of the Google Map showing semi-transparent polygons for each grouping of usage data, colored according to a predefined scale. The colour of each polygon indicates the level of wireless usage in that grouping, in this case 'building'. The figures are colour dependant and the scale range consists of nine values in descending RGB intensity.

Figure 4 depicted the node activity at a fixed point of time. In Figure 5 we plot the node location activity for a 24-hour period on 27 March 2009 (randomly selected). The campus 2D plane is converted into a 1D plan by aggregating all the campus node activity usage for a certain number of 1D chunks, in this case 14 chunks. The remaining 2 dimensions in Figure 5 show the number of associations and time (24 hours). These visualisations could be used, for example, by the network operators to look at the wireless network usage around the year and hence to find and predict usage patterns, congestion points, traffic and network growth. Note that it is possible for us to plot mobility patterns of each individual node or group of nodes across campus. It is also possible for us, for example, to find and plot relationships between individual nodes or group of nodes based on their activity patterns in time and space. Activity patterns may include mobility and presence in groups. However, collection, analysis and visualisation of such data may transgress on individuals privacy. We have looked at the analysis and visualisation of node location data with the aim to develop policies for the storage, usage, and release of node location data. The results presented in this section, for example, do not disclose any personal data of nodes, however these results are of limited use for network operators and could not be used to provide rich LBS's. Further discussion of the privacy issues and solutions is given in the next section.

5.2 Location Data and User Privacy – Swansea Considerations

One of the first criticisms of a Location Based Service (LBS) by its users is the perceived privacy implication of using a system where tracking user's movement is a fundamental principle of the technology. It is this alarming realisation that leads users to contemplate the possible effects that this data could have on their privacy, and in some cases causes an immediate rejection of the technology.

This reaction is understandable in a world where the terms *Big-Brother* and an *Orwellian State* are branded daily by the press, combined with headlines of governments on one hand wanting to create *Super Databases* [18] while with the other are losing personal data on buses and in the post [19]. Most users of networks in large institutions such as Universities are happy to use computers to browse the internet and communicate with friends and colleagues without considering how private that activity is. They are likely unaware that their internet browsing activity is being logged and the chat communications is being sent unencrypted and open to interception by network administrators or other agencies or hackers. It is only the fact that LBS immediately cause users to become consciously aware of the fact that their location data is what

drives the service and that data is being controlled by a computer somewhere that causes the kneejerk reaction. Privacy is considered a fundamental human right by the Universal Declaration of Human Rights and most democracies around the world [20]. Therefore the issue of location data and privacy must be taken seriously, and there has been substantial work to help provide privacy while also maintain the usefulness of LBS's.

One of the first considerations is the granularity of the location data. The granularity of meters could provide more information about a user than the granularity of kilometres [6]. The level of granularity in Swansea's LBS is per building, which removes a significant amount of detail from location data while still provide enough to offer a valuable LBS. Granularity alone does not provide any real privacy, and is vulnerable to correlation attacks as well as inference and assumptions based on historical data. One of the concerns of location data is how it is stored, and if the stored data could be abused by anyone who gains access, whether legitimately or not, to that data. One solution proposed to solve this is to anonymise the data using pseudonyms. Pseudonymity provides anonymity to location data while maintaining a relationship between the data that is used to help the LBS function. Recording a pseudonym and location as a location data record allows for the movement of a node to be tracked while removing any identifiable data from the record [21]. This adds a level of privacy to the system that would protect a user if the data was stolen or misused, however it does not offer complete privacy as a user's identity could still be inferred from the history of a nodes movement in some cases. One solution to this problem is the addition of dummy nodes that add a level of 'noise' to the LBS that does not affect the quality of the service but helps remove the ability of a possible attacker to infer the identity of a node based on the history of a nodes movements [22]. Another possible addition to add privacy is the use of mix zones which provide a trusted middleware that facilitates distribution of anonymised location information to third-party applications by defining spatiotemporal zones [23]. This does not directly fit into the system design at Swansea, but should be considered for the future.

The importance of privacy is fully understood by the research community and this is reflected in the amount of research undertaken in this area. But it has been suggested that the public put less significance on the importance and value of privacy and more on the short term benefits of the technology [24]. Regardless of the perceived value and importance of location privacy by the public it is vital that their privacy is maintained to the highest level, while still providing a valued service, in order to protect them from future and current threats to their human right of privacy.

6 Conclusions and Future Work

Systems security and user privacy have been major hurdles in the mass uptake of seamless mobile services. Location and context based approaches could provide additional network intelligence in securing networks. Privacy will have to be traded off for mobility and convenience; however, system designers could work to bring the trade-off equation more in the favour of privacy. We will continue to improve our work in these directions.

Acknowledgment. This work has been carried out partly through the JANET Location Awareness programme funds, whose support we acknowledge here.

References

1. Location Awareness Trial, Janet UK
2. Bahl, P., Padmanabhan, V.N.: RADAR: an in-building RF-based user location and tracking system (2000)
3. Harter, A., et al.: The Anatomy of a Context-Aware Application. In: Proc. 5th Annual ACM/IEEE Int. Conf. on Mobile Computing and Networking (1999)
4. Want, R., et al.: The Active Badge Location System. *ACM Transaction on Information Systems*, 1992 (10), 91–102 (1992)
5. Priyantha, N.B., Chakraborty, A., Balakrishnan, H.: The Cricket location-support system. In: Proceedings of the 6th annual international conference on Mobile computing and networking. ACM, Boston (2000)
6. Grlach, A., Heinemann, A., Terpstra, W.W.: Survey on location privacy in pervasive computing, in Privacy, Security and Trust within the Context of Pervasive Computing. The Kluwer International Series in Engineering and Computer Science. Kluwer Academic Publishers, Dordrecht (2004)
7. Hightower, J., Borriello, G.: A Survey and Taxonomy of Location Systems for Ubiquitous Computing, pp. 57–66 (2001)
8. González, M.C., Mehmood, R.: Experiences in Designing a UWB-based Indoor Localisation System. In: Eighth IASTED international conferences on Wireless and Optical Communications (WOC) 2008, Montreal, Canada. ACTA Press (2008)
9. Castro, P., et al.: A Probabilistic Room Location Service for Wireless Networked Environments. In: Proceedings of the 3rd international conference on Ubiquitous Computing. Springer, Atlanta (2001)
10. Jason Small, A.S., Seiwiorek, D.P.: Determining user location for context aware computing through the use of a wireless LAN infrastructure. *ACM Mobile Networks and Applications* 6 (2001)
11. Haeberlen, A., et al.: Practical robust localization over large-scale 802.11 wireless networks. In: Proceedings of the 10th annual international conference on Mobile computing and networking. ACM, Philadelphia (2004)
12. Youngjune, G., Jain, R., Kawahara, T.: Robust indoor location estimation of stationary and mobile users. In: INFOCOM 2004. Twenty-third Annual Joint Conference of the IEEE Computer and Communications Societies (2004)
13. John Krumm, J.C.P.: Minimizing Calibration Effort for an Indoor 802.11 Device Location Measurement System. Microsoft Research, MSR-TR-2003-82 (2003)
14. Koo, S.G.M., et al.: Location Discovery in Enterprise-based Wireless Networks: Implementation and Applications. In: Proceedings of the 2nd IEEE Workshop on Applications and Services in Wireless Networks (ASWN 2002), pp. 3–5 (2002)
15. Ferreres, A.I.G.T., Alvarez, B.R., Garnacho, A.R.: Guaranteeing the Authenticity of Location Information. *IEEE Pervasive Computing* 7(3), 72–80 (2008)
16. Aura, T., Roe, M., Murdoch, S.J.: Securing network location awareness with authenticated DHCP. In: Security and Privacy in Communications Networks and the Workshops. SecureComm. (2007)
17. Garg, S., Kappes, M., Mani, M.: Wireless access server for quality of service and location based access control in 802.11 networks. In: Proceedings of Seventh International Symposium on Computers and Communications. ISCC 2002 (2002)
18. Anderson, R., Brown, I., Dowty, T., Inglesant, P., Heath, W., Sasse, A.: Database State. Joseph Rowntree Reform Trust (2009)
19. Guardian, 25 Million Peoples Data Lost (2007)
20. Nations, U., Universal Declaration of Human Rights, General Assembly Resolution 217 A (III) (1948)

21. Chaum, D.: Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM* (24), 84–88 (1981)
22. Gruteser, M., Grunwald, D.: Anonymous Usage of Location-Based Services Through Spatial and Temporal Cloaking. In: *Proceedings of the 1st international conference on Mobile systems, applications and services*. ACM, San Francisco (2003)
23. Beresford, A.R., Frank Stajano, U.o.C.: Location Privacy in Pervasive Computing. *IEEE Pervasive Computing*, 10 (2003)
24. Acquisti, A., Grossklags, J.: Privacy and rationality in individual decision making. *IEEE Security & Privacy* 3(1), 26–33 (2005)