A Context-aware Trust-based Security System for Ad Hoc Networks

Maria Moloney, Stefan Weber
Distributed Systems Group
Department of Computer Science
Trinity College
Dublin 2, Ireland
mmolone2@cs.tcd.ie, Stefan.Weber@cs.tcd.ie

Abstract

Mobile ad-hoc networks (MANETs) comprise computer nodes which communicate over wireless links without any central control. Therefore, they must be able to make fully autonomous security decisions. This introduces new security challenges that existing security models and mechanisms do not adequately address [3]. In this paper we present a trust-based security system that deals with the specific challenges of MANETs by combining decentralised security management and context-aware computing. With this combination, our trust-based security system can establish appropriate trust levels for every situation.

1. Introduction

Security in computing is primarily concerned with achieving goals like authentication, confidentiality, integrity, anonymity, and availability. Traditional networks, such as infrastructure-based local area networks, use various security mechanisms to achieve these goals. However, with the exponential increase of mobile network devices, like laptops, PDAs and mobile phones, there is an increasing need for a change in the traditional architecture of networks. Fixed network elements are becoming unnecessary as mobile devices connect directly with one another and form networks on the fly or mobile ad-hoc networks (MANETs). This evolving architecture presents security challenges unaddressed by conventional security mechanisms [3]. Therefore, new security solutions need to be developed to maintain acceptable security levels throughout all network architectures.

In an ideal scenario, a MANET should self-organise and self-configure. All available devices on the network should work dynamically without any support from fixed infrastructure [12]. The network should be capable of carrying out routing and resource management while at the same time ensuring secure transmission of data.

In the next section, we outline the major security challenges of today's pervasive computing environment, which prevent this ideal scenario from becoming more widespread. We then outline two of our projects that address a number of issues related to MANETs. In particular, the first project, called SECURE, deals with trust in pervasive computing and the second one, called Aithne, focuses on sentient computing. In section three, we explain our security proposal which incorporates and builds upon various aspects of these projects. Section four explains our implementation proposal and finally section five outlines our conclusions and future work.

2. Trust-based Context-aware Security

In a traditional wired network every computer is physically secured to its environment by devices such as wires, alarms and locks. The digital security of these computers is controlled by traditional security mechanisms such as cryptography, firewalls and networks with dedicated administrative routers. These routers enforce digital security through a structured set of policies. A priori trust relationships between routers can be derived from these policies because the identity of each router is authenticated before admission to the network. This situation changes completely in pervasive computing environments. A key property of pervasive computing is that it contains mobile computers or devices. The devices within the environment are not physically secured and can move freely in and out of various mobile ad-hoc networks. Each mobile device has the potential to encounter thousands of other mobile devices within a short space of time. Therefore attempting to identify every device to enforce static security policies becomes impossible.

In addition, these devices are more vulnerable to physical security threats like theft as a result of being mobile. A party facing such a complex environment stands to benefit from interaction with these new devices, but only if it can assign meaningful privileges that allow mutual benefit while maintaining physical and digital security [2].

Our approach applies the human notion of trust to assign meaningful privileges to parties in pervasive

0-7803-9469-0/05/$20.00 ©2005 IEEE

computing environments. Trust naturally leads to a decentralized security management approach that can tolerate partial information, albeit one that has inherent risks for the trusting entity [2]. By clarifying the trust relationship between parties, logical and computational trust analysis and evaluation can be deployed. As a result, it becomes much easier to take proper security measures, and make correct decisions on any security issues [20]. Fundamentally, the ability to reason about trust and risk is what lets parties accept risk when interacting with other parties [2].

Our approach recognizes that trust is situation-specific; trust in one environment does not directly transfer to another environment [2]. A notion of context is particularly necessary when developing a trust-based security system for dynamic infrastructures like MANETs, where nodes frequently change environment [2].

The results of a questionnaire-based study in [10] showed that while the identity of the information requestor was a stronger determinant of privacy preferences than was the situation in which the information was collected, situation was nonetheless an important determinant.

Context has been defined in various ways. Some definitions are broader than others. Schilit and Theimer [13] define context as only ‘‘location and the identity of nearby people and objects’’. Ryan et al. [12] broaden the meaning to encompass ‘‘location, identity, environment, and time’’. Dey [5] goes further and defines context as ‘‘any information that can be used to characterise the situation of entities’’ which is ‘‘typically the location, identity and state of people, groups, and computational and physical objects.’’. Schilit et al. [14] see context to mean:

Context encompasses more than just the user's location, because other things of interest are also mobile and changing. Context includes lighting, noise level, network connectivity, communication costs, communication bandwidth and even the social situation, e.g., whether you are with your manager or with a co-worker.

The last and broadest definition by Schilit et al allows us to exploit physical location and any other information about users and resources to enhance the user experience. [3].

Traditionally, security requirements did not need to be context-sensitive as computing existed within a static environment. However, as computing technology becomes more and more integrated into everyday life, it is essential that security mechanisms become more flexible and less intrusive [4]. Security in pervasive computing must be able to assimilate changes in context and situational information effortlessly [3]. For example, access control decisions should factor in time or location and they should be able to change dynamically to limit permissions to times or situations when they are needed. However, viewing what the security policy might become in a particular time or under a particular situation should not be possible [3]

Existing trust-based security solutions fail to take into account the dynamic and situation-specific nature of pervasive computing environments and more specifically of MANETs. [9] Additionally, the majority of existing proposed trust-based solutions for MANETS have been defined through formal methods and only a subset of these have been evaluated through simulations. Typically, simulators attempt to replicate real-world scenarios; but their results are only as good as the initial data and configurations. It has been shown that these results can diverge significantly from the experience in real-world scenarios [7].

In this paper, we propose a new type of trust-based, context-aware security system specifically for MANETs. We intend to evaluate our system through a series of combined simulations and real-world tests. The results of each series of tests will go towards refining and redefining the design and specification of our solution to eventually provide a complete solution specifically for the MANET environment.

To help us with the initial stages of our project, we sought the expertise of two existing projects in Trinity College Dublin.

2.1. SECURE

The aim of the SECURE project is to produce an advanced, formally grounded and reusable trust-based security framework (TSF) [15]. This TSF is designed to function on any network infrastructure. However, to date, it has not been tested or evaluated for MANETs.

Figure 1 shows the basic components of the TSF. It depicts a decision-making component that is called when a requested entity has to decide what action to take upon a request made by another entity, the requesting entity.

In order to make this decision, two sub-components are used:

- a trust engine that can dynamically assess the trustworthiness of the requesting entity based on pieces of evidence e.g. observations or recommendations [19].
- a risk engine that can dynamically evaluate the risk involved in the interaction and choose the action that would maintain the appropriate cost/benefit.

In the background, another component, the Evidence Manager, is in charge of gathering evidence. The available evidence can be drawn from records of previous interactions or recommendations from partly trusted third parties [2].

This evidence is used to update both risk and trust information. Thus, trust and risk follow an interdependent, managed life-cycle.

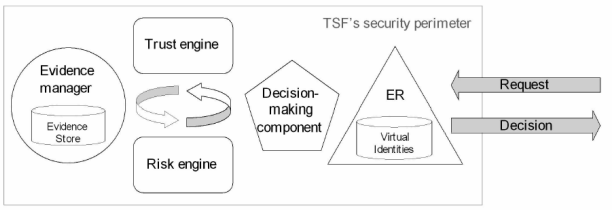

Figure 1 High-level view of a Trust-based Security Framework [14].

The Entity Recognition [15] module is the part of SECURE that incorporates context information into the decision making process. SECURE follows Dey's definition of context which is typically:

"the location, identity and state of people, groups, and computational and physical objects." [5]

From this list, SECURE sees identity recognition as the only important context information for computing trust. It incorporates this into its TSF through its Entity Recognition Module. It does not, however, handle any other type of context information other than identity. This means information regarding, among other things, a party's location and activity are disregarded when making security decisions in SECURE.

We propose, for our security system, to devise an instance of the SECURE framework specifically adapted for MANETs which processes other context-aware information through its evidence manager and adds it to its virtual identities stored in its Entity Recognition Module. In this way, the virtual identities will have different trust values depending on their context at any given time. Figure 2 shows a top level view of our adapted instance of the SECURE TSF.

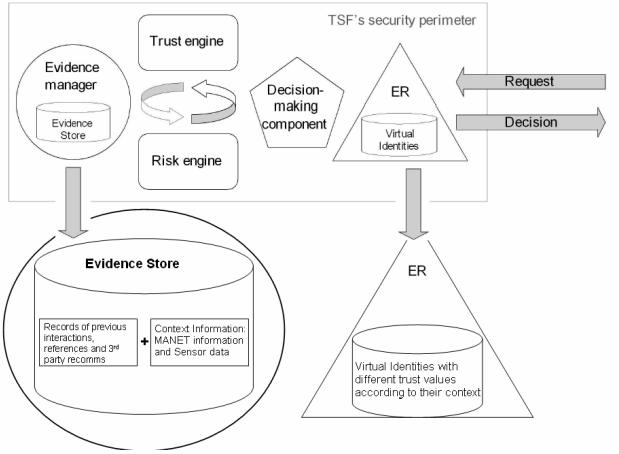

Figure 2 The Adapted Instance of the SECURE TSF

This instance improves the calculation of trust in the SECURE framework, which in turn will strengthen it as a complete security solution for MANETs. We aim to achieve this context-aware solution with the help of a second project in Trinity College Dublin.

2.2. Aithne

The Aithne project aims to design, implement and evaluate a middleware architecture to support application development for sentient computing in areas ranging from environmental monitoring and control to Intelligent Transportation Systems (ITS). The middleware architecture will support applications that handle large collections of collaborating sensor-rich computational devices.

The programming model used in the Aithne project is based on sentient objects. The latter are software components that lie in the control path between at least one sensor and one actuator. They can both consume and produce events.

A sensor is seen as an entity that produces software events in reaction to a real-world stimulus detected by some physical device, whereas an actuator is defined as a component that consumes software events and reacts by attempting to change the state of the real world in some way via some physical device [17].

Figure 3 shows a sentient object and its internal workings. Essentially, sentient objects sense and interact with their environment via sensors and actuators. It is this awareness of, and interaction with the environment that makes context awareness an important factor in sentient objects. Aithne uses a very broad definition of context. It sees context as:

any information sensed from the environment that may be used to describe the situation of a sentient object. This includes information about the underlying infrastructure available to the sentient object [17].

and context-awareness as:

The use of context to provide information, to a sentient object. This information may be used in its interactions with other sentient objects, and/or the fulfillment of its goals [17].

As outlined in Figure 3, the sentient object programming model supports context acquisition (sensory capture), context representation and inference [17].

2.2.1. Context Acquisition

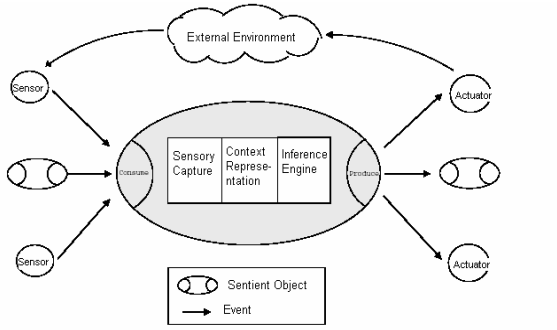

Figure 3 A Sentient Object from Aithne [17].

A sentient object may receive input from an array of diverse sensors, for example a sentient vehicle's array of sensors could include proximity sensors, GPS, speed and direction sensors, and pollution sensors. Signals from these sensors need to be integrated in order to determine the overall environment and context of the sentient object [17].

The type of context acquisition that will occur in our security solution will capture context-information which will help to assess the security risks in the current environment. Our context acquisition component will be adapted to collect information on both physical security risks in the environment and digital security risks of the MANETs in the environment.

2.2.2. Context Representation

Raw sensor data will usually need to be transformed in some way before it may be considered useful contextual information. Such transformation may occur in the sensor itself, or may be carried out within the sentient object itself. The context representation component deals with the representation of context information in a way that is useful to the sentient object and may be easily exchanged amongst sentient objects [17]. Our context-representation component will represent security related context information for our system.

2.2.3. Inference

Sentient objects are expected to act upon the information they receive from their environment and change its state. This process implies some form of decision making ability or intelligence, on the part of the

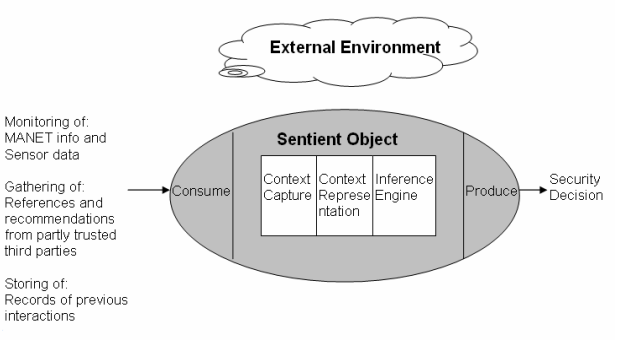

Figure 4 Our adapted Sentient Object

sentient object that is captured in the inference engine component. An inference engine, in artificial intelligence refers to a program that reasons about a set of rules (a knowledge base) in order to derive an output. The knowledge base of an inference engine contains knowledge required to solve a certain problem, encoded as a set of production rules. The knowledge encoded in rules is generally captured from a human expert who expresses his expertise in the form of such rules. Figure 4 shows our adapted sentient object.

The aim of our inference engine is to act upon security information and implement the appropriate security levels. Our inference engine will host our adapted instance of the SECURE TFS. We will combine the ability of SECURE to gather evidence through recommendations and references with the inference engine of our sentient object which is capable of reasoning about context. This combination will be stored in the evidence store of the adapted instance of SECURE. Then, once a request for interaction occurs our instance of SECURE will assess the trust and risk levels by using the evidence store of references, recommendations and context information. Figure 5 outlines the final structure of our security solution.

The inference engine developed in the Aithne project is as generic as possible so that it may be applied to a number of different knowledge bases in different domains with minimal changes to itself [17]. For the purposes of SECURE, we can envisage a sentient object inference engine which incorporates the context-aware instance of the SECURE TSF we are using for our security solution.

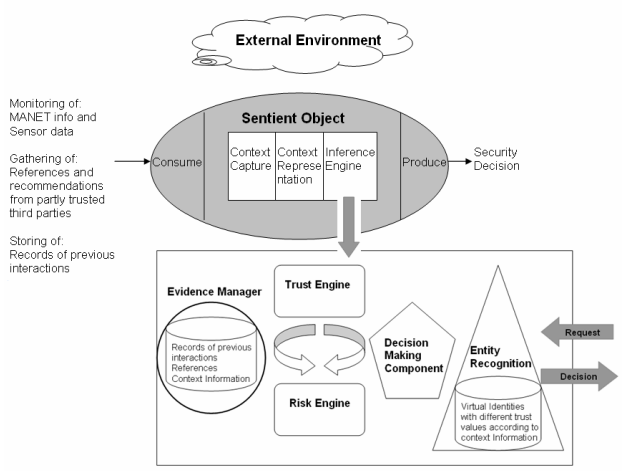

Figure 5 Our Context-aware Security System.

3. The Innovation

Initially, we intend to use the original TSF supplied by the SECURE project to examine and highlight problem areas specific to MANETs, in various existing trust models. By rigorously evaluating these trust models through simulations and real-world scenarios, we hope to gain a better understanding of the adequate provision of trust-based security for MANETs.

The results of these evaluations will help us to devise an instance of the SECURE TSF with context-aware capability which identifies the inherent security risks specific to MANETs and makes appropriate trust-based security decisions accordingly.

By processing information about the environment, nodes in a MANET can employ different levels of trust according to the security requirements of their current environment.

When a node enters a new MANET, it can attempt to identify security risks by gathering and processing context information from three main sources: the first source is sensor data from the environment. This can be information such as whether the environment is hot or cold, or whether it is day or night. It could also be information gathered from other nodes on the MANET regarding their location and activity on the MANET.

The second source of context information is MANET information such as hop counts and network topology. This information can help assess the reliability of received messages. For example, a message from a node that is two hops away is more reliable than a message from a node that is three hops away. This is due to the circumstance that, the less the hop count, the less likelihood of tampering

The third source of context information consists of identity information in the form of records of previous transactions and references which consist of recommendations from partly trusted third parties. The SECURE TSF has the capability of gathering, storing and reasoning about 'identity' context information. These three sources of information are processed in the evidence store of our system. Figures 2 and 5 outline the information stored in the evidence store of our system.

Given the new context information, a node can adjust its own security levels for each new MANET it joins and also update its existing trust relationships with nodes on familiar MANETs.

3.1. A Sample Scenario

The working of our trust-based, context-aware security system can be more clearly explained through a simple scenario example. Suppose Jim enters a street where a MANET has been established. His PDA detects the new MANET and gathers specific information about nodes on the MANET to establish basic trust relationships with them. Initially, the trust between Jim's PDA and the other nodes is low.

Jim's PDA is informed by two different "Coffee shop" nodes on the MANET that Jim can place an order for a coffee with them and collect it as he passes further down the street.

Because Jim is new to the MANET the reputation of both shops is unknown to him. Therefore, Jim's PDA sets about finding out, in six steps, which shop offers the best coffee.

Step 1: Initially, Jim's PDA gathers sensor data from the street and the environment. It gathers information such as the length, width and cleanness of the street. It examines the time of day Jim is entering the street and whether it is day or night. It investigates whether the street contains mobile, vehicle or business nodes in order to determine whether the street is a busy shopping street or a quiet, back street. With the gathered context information, it decides that the street is a safe area for Jim to enter because it is a busy shopping street during the morning of a weekday.

It is important to note here that if it the environment and therefore the context information being gathered were different, for example, it was the middle of a Saturday night and the street was full of other mobile nodes but no business nodes, Jim's PDA may decide the street was not safe to enter.

Table 1. Sensor data gathered from the current environment.

LOCATION	TIME	ENTITY	ACTIVITY
Static or mobile	DD/MM/YY	Business/ sole trader	Coffee/pizza

Step 2: It then gathers sensor data from each of the 'coffee shop' nodes as shown in Table 1. Jim's PDA looks at both of their locations and sees if they are either static or mobile. Static would indicate that the shop is located in a building and mobile would indicate that the shop is in fact a street stall. The time value would indicate how long the shop has been in business and the entity value would specify whether they are a business or something else like a sole trader. The activity value would show whether the shops specialise in coffee.

Step 3: Jim's PDA seeks further recommendations about the coffee shops from other nodes on the network. It gets both positive and negative feedback/references from the neighbouring nodes. However, it does not consider each reference as equal. It rates them according to the sensor data it gathers from each neighbour. It gathers the same data, as outlined in table 1 above (location, time, entity and activity), from neighbours as it does from the two coffee shops. With this information, it assesses the trustworthiness of each reference. For example, other coffee shop references, which are unfavourable, are disregarded because they may deliberately give bad feedback for selfish purposes. Businesses on the street are regarded highly as they are probably frequent customers of the coffee shops. Mobile nodes on the street are also considered since, if they have references for the coffee shops in their Evidence Manager, then they have had previous interactions with the shops.

Step 4: Jim's PDA gathers MANET information. The sensor data gathered from the 'coffee shop' nodes and the referees is assessed according to the amount of hops they are from Jim's PDA.

Step 5: finally, Jim's PDA must calculate, with the aid of the sensor data and MANET information gathered from both coffee shops and referees, combined with the appraised references received from the neighbouring nodes, which coffee shop is a better choice. Once a choice has been calculated, the PDA recommends the most trustworthy choice. It is up to Jim to decide to act upon the recommendation.

Step 6: Once the transaction is finalized Jim either voluntarily inputs whether the calculated choice was a good one or his PDA actively requests feedback regarding the transaction. It then stores this information, for future reference, in its evidence store as a "record of previous interactions" for the recommended coffee shop. It stores all the gathered context information which influenced the transaction such as the condition of the street, time of day etc.

Figure 6 shows how all this information is processed so that Jim's PDA can derive a decision and make a recommendation to Jim about which coffee shop he should choose.

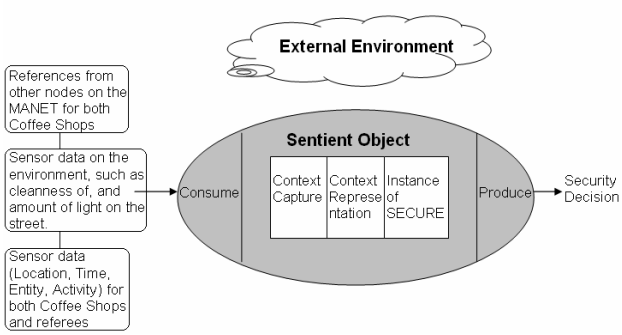

Figure 6 Information processing in Jim's PDA

4. Implementation

The sole purpose of this position paper is to stimulate discussion of our security system in a workshop on security through collaboration. This section gives a broad outline of the proposed implementation process of our security system.

The Distributed Systems Group in the Department of Computer Science at Trinity College Dublin, has deployed the Wireless Ad hoc Network for Dublin (WAND) as a large-scale test-bed for MANET protocols and applications to provide an opportunity to explore the behavior and performance of a variety of routing protocols in a real life environment, and to investigate the use of mobile applications in an urban environment [1].

We intend to use this test-bed to determine the advantages and disadvantages of our security system and establish the exact circumstances under which our system represents a good choice.

Initially, we will evaluate the suitability of the original SECURE TSF through laboratory simulations and through real-world testing on WAND. This will provide a baseline against which we can determine the effectiveness of our adapted instance of SECURE for MANETs. In addition, it will facilitate the process of assessing and making recommendations for improving our instance of the SECURE TSF.

The first evaluation lifecycle will entail a thorough review of the SECURE TSF architecture and configurations. The review will involve meeting with members of the SECURE project and reading related documentation to gain an understanding of the framework before installing it on simulated MANETs and our real-world MANET, WAND. Once SECURE is installed and configured on the MANET a series of penetration tests will be performed and the results recorded.

Penetration tests are security tests in which evaluators attempt to circumvent the security features of a system based on their understanding of the system design and

implementation. In this scenario, we will attempt to circumvent the SECURE TSF in order to carry out attacks on the network.

Once all the necessary tests are complete, the information gathered during the first evaluation lifecycle will provide a basis for understanding the suitability of the SECURE TSF for MANETs. We will analyze the security weaknesses that were highlighted during the evaluations and adapt the design of our instance of the SECURE TSF to incorporate countermeasures to combat these weaknesses.

The next step is a second evaluation lifecycle to assess our adapted instance of the SECURE TSF using the same tests under the same conditions as the first lifecycle.

Once the second lifecycle is complete, we can compare the results of both evaluations and assess the efficacy of our new security system.

The WAND test-bed offers us a unique opportunity to explore the behaviour and performance of existing trust models and to develop a new security system in a real-life environment that reflects all the randomness and unpredictability that is extremely difficult to reproduce with simulation.

5. Conclusion

This paper outlines our plan to provide a context-aware trust-based security system for MANETs. It incorporates a sentient object model within which resides a trust-based security framework adapted specifically for MANETs. This combination of context-awareness and trust reasoning allows our system to calculate the trust value of an entity based on previous interactions with the entity, references from partly trusted third parties and sensor data gathered from the current environment.

If successful, the new system will improve the calculation of trust for mobile applications, which in turn will provide a stronger and more complete security solution specifically designed to combat the inherent security weaknesses of MANETs. Finally, we intend to test our system on a real world test bed provided by the WAND project in Trinity College.

6. References

[1] P. Barron, S. Weber, S. Clarke, and V. Cahill, "Experiences Deploying an Ad-hoc Network in an Urban Environment," presented at REALMAN: IEEE ICPS Workshop on Multi-hop Ad hoc Networks: from theory to reality, Santorini, Greece, 2005.

[2] V. Cahill, E. Grey, J.-M. Seigneur, C. Jensen, and Y. Chen, "Using Trust for Secure Collaboration in Uncertain Environments," in *IEEE Pervasive Computing Magazine*, vol. 2, 2003.

[3] R. Campbell, J. Al-Muhtadi, P. Naldurg, G. Sampemane, and M. D. Mickunas, "Towards Security and Privacy for Pervasive Computing," presented at The International Symposium on Software Security, Keio University, Tokyo, Japan, 2002.

[4] M. J. Covingtony, P. Fogla, Z. Zhan, and M. Ahamad, "A Context-Aware Security Architecture for Emerging Applications," presented at 18th Annual Computer Security Applications Conference, Las Vegas, Nevada, 2002.

[5] A. K. Dey, "Understanding and Using Context," *Personal and Ubiquitous Computing Journal*, vol. 5, pp. 4-7, 2001.

[6] A. Fitzpatrick, G. Biegel, S. Clarke, and V. Cahill, "Towards a Sentient Object Model," presented at Workshop on Engineering Context-Aware Object Oriented Systems and Environments(ECOOSE), Seattle WA, USA., 2002.

[7] G. Gaertner, E. ONuallain, A. Butterly, K. Singh, and V. Cahill, "802.11 Link Quality and its Prediction - An Experimental Study," presented at 9th Intl. Conference on Personal Wireless Communications (PWC 2004), Delft, The Nederlands, 2004.

[8] S. Giordano, "Mobile Ad-hoc Networks," in *Handbook of Wireless Networks and Mobile Computing*: John Wiley & Sons, pp. 325-346, 2000.

[9] E. Gray, P. O'Connell, C. Jensen, S. Weber, J.-M. Seigneur, and C. Yong, "Towards a Framework for Assessing Trust-Based Admission Control in Collaborative Ad Hoc Applications," Department of Computer Science, Trinity College, Dublin TCD-CS-2002-66, December 2002.

[10] S. Lederer, J. Mankoff, and A. K. Dey, "Who Wants to Know What When? Privacy Preferences Determinants in Ubiquitous Computing.," presented at The ACM Conference on Human Factors in Computing Systems, Fort Lauderdale, Florida, 2003.

[11] C. S. R. Murthy and B. S. Manoj, *Ad Hoc Wireless Networks Architectures and Protocols*. New Jersey: , Prentice Hall PTR, 2004.

[12] N. Ryan, J. Pascoe, and D. Morse, "Enhanced reality Fieldwork: the context-aware archeological assistant," *Gaffney, Leusen and Exxon (eds) Computer applications in archeology*, 1997.

[13] B. Schilit and M. Theimer, "Disseminating active map information to mobile hosts," in *IEEE Network*, vol. 8, pp. 22-32, 1994.

[14] B. Schilit, N. Adams, and R. Want, "Context-aware computing applications," presented at Proceedings of the IEEE Workshop on Mobile Computing Systems and Applications, Santa Cruz, CA, 1994.

[15] J.-M. Seigneur and C. D. Jensen, "The Role of Identity in Pervasive Computational Trust," presented at The Second International Conference on Pervasive Computing, Vienna, Austria, 2004.

[16] J.-M. Seigneur and C. D. Jensen, "The Claim Tool Kit for Ad-hoc Recognition of Peer Entities," *Journal of Science of Computer Programming*, vol. 54, pp. 49-71, 2005.

[17] A. Senart, "Aithne: Middleware for Sentient Computing," vol. 2005: Trinity College Dublin, 2004.

[18] A. Tanenbaum, *Computer Networks*, 4th ed: Upper Saddle River, N.J. :Prentice Hall., 2002.

[19] W. Wagella, S. Terzis, and C. English, "Trust-Based Model for Privacy Control in Context-Aware Systems," presented at Proceedings of the 2nd Workshop on Security in Ubiquitous Computing, Seattle, Washington, USA, 2003.