Trust Propagation in Small Worlds

Elizabeth Gray[1], Jean-Marc Seigneur[1], Yong Chen[1], and Christian Jensen[2]

[1] Distributed Systems Group, Department of Computer Science
Trinity College, Dublin 2, Ireland
{grayl, seigneuj, cheny}@tcd.ie
[2] Informatics and Mathematical Modelling, Technical University of Denmark
Richard Petersens Plads, Building 322, DK-2800 Kgs. Lyngby, Denmark
cdj@imm.dtu.dk

Abstract. The possibility of a massive, networked infrastructure of diverse entities partaking in collaborative applications with each other increases more and more with the proliferation of mobile devices and the development of ad hoc networking technologies. In this context, traditional security measures do not scale well. We aim to develop trust-based security mechanisms using small world concepts to optimise formation and propagation of trust amongst entities in these vast networks. In this regard, we surmise that in a very large mobile ad hoc network, trust, risk, and recommendations can be propagated through relatively short paths connecting entities. Our work describes the design of trust-formation and risk-assessment systems, as well as that of an entity recognition scheme, within the context of the small world network topology.

1 Introduction

The proliferation of mobile devices and development of vast ad hoc networks introduces the possibility of an environment where multitudes of diverse entities will partake in collaborative applications with each other. A mobile ad hoc network is an autonomous system of mobile entities connected by wireless links. All entities are free to move randomly, and the network is self-organising, which makes it highly dynamic and subject to rapid and unpredictable changes. As in traditional networks, access to collaborative resources in mobile ad hoc networks requires varying levels of control. Also, some way of authenticating an entity is needed, as well as a way of determining what access that entity may have to shared resources. Traditional authentication and access control methods fail when applied in a decentralised collaborative ad hoc environment. For example, in traditional groupware applications, access to a group is controlled by an administrator with a predefined list of names and access permissions of group members. The administrator grants access rights based on whether the requesting entity is authenticated and identified as meeting the appropriate criteria. However, in a network that is constantly changing both size and topology, this approach does not scale.

This is best illustrated by the following example. Suppose that while on the 8am commuter train every weekday, Alice joins an ad hoc wireless network to see

what collaborative gaming applications are available. One morning, she discovers a blackjack game in which Bob is the dealer, and she requests admission to the game. To Bob, Alice is an unknown entity, who may or may not be trusted to behave correctly, e.g. pay her gaming debts, if given access to his game. In the traditional model, Bob would be able to contact a centralised administrator to determine if Alice is Alice, and if she should have access rights to participate in the blackjack game. This example shows that traditional authentication methods do not scale to the large mobile ad hoc networks envisioned.

We propose a solution for this scenario based on the human notions of trust, risk, and recognition in human ad hoc collaborative networks. Every day, humans determine how to interact with known, partially-known, and unknown people. Much of the time, we do this with no assistance from a trusted, centralised third party and without the availability of complete information. Humans use the concepts of trust, risk, and recognition to help decide the extent to which they cooperate with others. In this way, mechanisms are provided for lowering access barriers and enabling complex transactions between groups.

Difficulties lie in trying to map the human concepts, which themselves are defined differently across the various fields of research, to a computational model. A first important step in this domain is given by Marsh [1], who demonstrates that the concept of human trust can be formalised as a computational model. Another critical step comes from McKnight and Chervany [2], who describe a framework for regulating trust formation, so that unambiguous conversation between computational entities can occur. We implemented McKnight and Chervany's trust framework [3], whereby a trust-based admission control system allows entities in open, diverse systems a way of directly establishing trust in one another. Within this system, as trust formation occurs, trust is measured and used to make dynamic admission control decisions. A problem arises the first time two previously-unknown entities interact, because Alice has to decide her initial trust in Bob.

Further comparisons with human networking concepts give us a possible solution to this problem, based on recommendations from mutual acquaintances. Sociologists estimate that each human has roughly 300 acquaintances with whom he is on a first-name basis. This means that there are 300 people one step away from any given person, 90,000 people two steps away, 27 million people three steps away, etc. This sociological concept is the basis for small world research [4], which describes the tendency for each entity in a large system to be separated from any other entity in the system by only a few steps. Small world research formalises human networking concepts, and gives us standard formulae with which to analyse seemingly random digital networks. We aim to describe how trust-based security measures can be furthered by developing a design against the backdrop of small world theories.

In this paper, we describe how, within the context of the small world network topology, the human concepts of trust, risk, and recognition can be applied to secure collaborative applications in mobile ad hoc networks. The structure of the paper is as follows. Section 2 is an examination of the small world theory.

Section 3 specifies the design of our trust-based security architecture within the context of small worlds. Finally, Section 4 presents conclusions and ideas for future work.

2 Small Worlds

The small world concept suggests that any pair of entities in a seemingly vast, random network can actually connect in a predictable way through relatively short paths of mutual acquaintances. The work on small world theory is of significant interest to our research on the formation and propagation of a computational notion of trust within collaborative networks that appear to be made up of completely random and dynamic connections.

In this section, we review the major research into small world theory, including Stanley Milgram's seminal work in the field as well as more recent developments from the areas of sociology, psychology, and computer science.

2.1 Small World Beginnings

In the 1960's, Stanley Milgram [5], a social psychologist at Harvard, researched the hypothesis that members of any large social network are connected to each other through short chains of intermediate acquaintances, as described in the handshaking scenario above. Milgram and Travers [6] brought small world theory to the attention of the academic community in 1969, having performed the following experiment to prove the hypothesis. They sent information packets to a few hundred randomly selected people in Nebraska and Kansas. Within each information packet was the name of one of two target persons in Boston. Each Nebraskan/Kansan was to forward the information packet onto some acquaintance known on a first-name basis, until the packet reached the target person in Boston. The famous result is one of 'six degrees of separation,' which states that any two people in the U.S. population at the time are connected by no more than six steps.[3]

This result is startling because a person's average number of acquaintances is significantly less than the size of the entire population. Milgram's results are even more surprising when we consider that a person's acquaintances are not generally spread evenly throughout the population. Instead, acquaintanceship tends to be based on common location, background, interests, etc. Therefore, most of a person's acquaintances would be in a tight network, or clique, around him. Within each clique is a high level of redundancy, i.e. within Alice's circle of friends, most of them are also friends with each other. In theory, if all human networks are based upon tight, closed, cliques, it would take far more than a few steps to link two strangers in populations of millions.

[3] The Department of Sociology at Columbia University is currently carrying out the first large-scale global verification of Milgram's small world hypothesis, using email rather than the postal service.

Watts and Strogatz [7] furthered Milgram's ideas by modelling small world networks as distinguished from ordered networks and random networks, as shown in Figure 1. Assuming that any network can be represented by connections existing between its members, broad classes of networks can be defined with a range between highly ordered and highly random. In the fully ordered case, the network is completely regular and cliquish. One node knows only the nodes immediately adjacent. In this type of network, many steps are required to connect non-adjacent nodes. The second case is a totally random network, wherein no cliquish behaviour is exhibited. In this type of network, a node is just as likely to be linked to an adjacent node as to a non-adjacent node. The intermediate case that Watts and Strogatz were able to model is the small world network. Randomness is introduced into a fully ordered network by randomly adding 'shortcuts', which are links from one point to another point in the network that would usually take several steps to access. In a small world, any given node has an immediate clique of adjacent connections and may or may not also be connected via a shortcut to a node in any other part of the network. In fact, just a very small number of random links is enough to 'short circuit' an otherwise huge, ordered network. For example, if only one node in 100 has a random link to any other node in the network, the average number of steps linking network node pairs decreases tenfold. Therefore, a small world network has the characteristics of a fully ordered network, but as randomness increases, the number of steps needed to link nodes decreases.

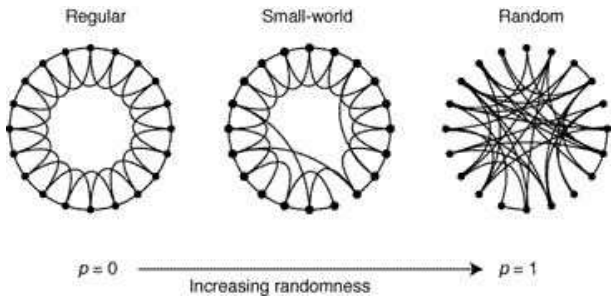

Fig. 1. Increasing Network Randomness [7]

To give a better understanding of the small world network model, Watts [8] identifies three characterising properties. The first is characteristic path length, which is the shortest path required to connect one node to another. This is averaged over all node pairs to give the characteristic path length of the whole network. The second parameter is the clustering coefficient, which measures the probability that two nodes that are connected via a mutual acquaintance will also be directly connected to one another, i.e. the cliquishness of the network. Watts shows that, according to these two parameters, highly ordered networks have long characteristic path lengths and large clustering coefficients, while random networks have short characteristic path lengths and very little clustering. A

small world network exhibits characteristic path lengths approximately as short as those of random networks, but with much greater clustering. The third property is logarithmic length scaling such that for all graph sizes, as the network graph grows significantly larger, the average characteristic path length remains relatively small.

Watts and Strogatz assess three real-world networks: social, power, and neural. Each of these three networks exhibits small world topologies. For example, the connections exhibited by the social network [9] are delineated by short paths between a given actor (most notably, actor Kevin Bacon in the 'Kevin Bacon Game' developed at the University of Virginia [10]) and any other actor in the population of actors. The Kevin Bacon Game small world network exhibits significant clustering and a small characteristic path length that remains relatively small, no matter how large the database of actors grows.

Adamic [11] extends Watts' research to prove that another real-world network, the World Wide Web, is also small world. Based on a sample of .edu sites, at the site level the Web exhibits an average characteristic path length of four and a clustering coefficient significantly higher than in a random network of similar size.

2.2 Identity and Search in Small Worlds

While models such as that of Watts and Strogatz present excellent analysis of Milgram's conclusions regarding the pervasiveness of short chains in a range of real-world networks, Kleinberg [12] finds these models insufficient to explain a second component of Milgram's findings: 'that individuals using local information are collectively very effective at actually constructing short paths between two points in a social network.' Kleinberg extends Milgram's original research to illustrate that not only is it possible for networks to have short characteristic path length and local clustering, but also that it is possible for an entity to use local information to find short paths without requiring a map of the entire network.

Kleinberg defines an infinite family of random network models that generalizes the Watts-Strogatz model. For one of these models, then, he shows that there is a decentralized algorithm capable of finding short paths with high probability. Finally, he proves that there is a unique model within that family for which decentralized algorithms are effective for navigation.

Kleinberg specifically focuses on decentralised algorithms, i.e. those by which is passed sequentially from an entity to one of its local or long-range connections using only local information. It is stressed that constraining the algorithm to use only local information is crucial to this research because if an entity had knowledge of all other entities in the network, it could simply perform a breadth-first search to locate the shortest path.

Watts et al [13] incorporate similar ideas into their research of social networks. They define the concept of 'searchability,' the property of being able to find a target quickly in a networks. The model gives an explanation of social networks in terms of searchability based on recognizable personal identities,

where identity is considered to be a set of characteristics measured along social dimensions. A class of searchable networks is defined, as is a method for searching which, similar to that of Kleinberg, is a decentralised algorithm based on Milgram's work whereby each entity forwards a message to its neighbour who is closer to the target entity in terms of social distance. This research suggests that searchability is a generic property of real-world social networks, and that an effective decentralised search can be conducted provided that two pieces of information are known: the characteristics of the target entity, and the current entity's immediate neighbours.

In the case of the search algorithms presented by Kleinberg and Watts, there is still the underlying assumption that each entity in a network has a findable, unchanging location. This assumption does not hold in mobile ad hoc networks, which do not rely on any fixed infrastructure. In this type of network, all networking functions must be performed by entities themselves in self-organising manner. In this regard, Hubaux et al [14] present their Shortcut Hunter algorithm, which shows that certificate chains result with high probability between two previously unknown entities using only their merged local certificate repositories. Capkun et al [15] build on Hubaux's work and propose a new approach to securing mobile ad hoc networks. The work is PGP-based, as PGP's functionality relies solely on user acquaintances, and shows that the small world phenomenon naturally emerges in the PGP system as a consequence of the self-organisation of the users. Moreover, Capkun et al argue that self-organised security systems in which entities issue certificates based on acquaintanceship will exhibit small world properties as a result of the formation of mutual trust relationships.

2.3 Summary

In this review, we find that self-organising networks, such as particular types of mobile ad hoc networks, exhibit small world tendencies. This means that we may use existing distributed algorithms, developed in small world research, to establish shortcuts between cliques in mobile ad hoc networks. Based on this premise, we see that small world characteristics become increasingly relevant to the design of a security system that incorporates the elements of trust, risk, and entity recognition. In the following section, we present our design concepts for such a system.

3 Trust-Based Security Mechanisms in Small Worlds

In this section, we present a trust-based security architecture, including the design of four components that may be used to provide security in mobile ad hoc networks: entity recognition, trust-based admission control, risk assessment, and trust management. Each component's design is heavily influenced by the concepts illustrated in small world research.

3.1 Trust-Based Security Architecture

An overview of our trust-based security framework is shown in Figure 2.

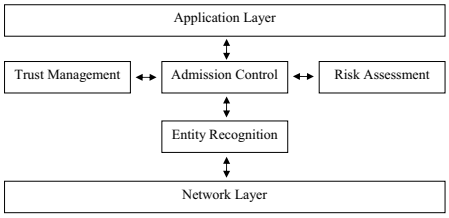

Fig. 2. Trust-Based Security Architecture

The framework consists of four main components. *Entity Recognition* observes encountered entities and decides whether they have been encountered before. *Trust-Based Admission Control (TBAC)* examines the recognised entity and decides whether sufficient trust exists to offset the risk involved with collaboration with this entity. *Risk Assessment* examines the recognised entity and calculates the risk involved with collaboration with this entity. *Trust Management* manages the recorded experiences from interactions with encountered entities.

We return to the collaborative gaming example from Section 1 in order to illustrate how these components are used. Alice wishes to join a blackjack game in which Bob is the dealer. She sends a message to Bob requesting to join the game. The entity recognition component on Bob's device determines whether Bob knows Alice from previous interactions, i.e., Bob tries to *recognize* Alice. The result from the recognition process (either "Alice" or "unknown") is passed on to the TBAC component which decides whether Alice is allowed to join or not. In order to perform its task, the trust-based security framework passes the result from the recognition process to the risk assessment component in order to determine the risk involved with interacting with Alice. Once the risk has been established, the TBAC component queries the trust management component in order to determine the level of trust that has been established from previous interactions with Alice. If the level of trust is sufficient to offset the risk, Alice is allowed to join the game. Bob's TBAC component may also consult the current players, but this protocol is beyond the scope of this paper. The four components are discussed in greater detail in the following sections.

3.2 Entity Recognition

Recognition is a notion humans use when interacting with one another. A person recognises, partially recognises, or does not recognise another person, and this helps him determine the level of trust in the other person and to assess the risk of a particular interaction.

A similar concept arises in mobile ad hoc networks, where diverse entities must interact in a highly dynamic and unpredictable environments. Traditionally, authentication is the first block to ensure secure computing [16], because, without being certain with whom an entity interacts, three fundamental properties - confidentiality, integrity and availability - can be trivially violated through interaction. Authentication, however, does not scale to the world of ubiquitous computing. We believe that, in this context, it is more beneficial to take an approach based on entity recognition [17, 18], rather than solely on traditional authentication schemes such as PKI [19] or Kerberos [20].

Authentication Process (AP)	**Entity Recognition (ER)**
A.1. **Enrolment**: generally involves an administrator	
A.2. **Triggering**: e.g., someone clicks on a Web link to a resource that requires authentication to be downloaded	E.1. **Triggering** (passive and active sense): mainly triggering (as in A.2), with the idea that the recognizing entity can trigger itself
A.3. **Detective work**: verification of the principal's claimed identity	E.2. **Detective work**: to recognize the entity using the negotiated and available recognition scheme(s)
	E.3. **Retention** (optional): whereby there is preservation of after-effects of experience and learning to make recognition possible
A.4. **Action**: the identification is subsequently used in some way. The claim of the identity may be done in steps 2 or 3 depending on the authentication solution. (Loop to A.2.)	E.4. **Action** (optional): the outcome of the recognition is subsequently used in some way (Loop to E.1.)

Fig. 3. Comparison of Authentication and Entity Recognition[18]

In general, authentication schemes start with enrolment of entities. This is a static process, which allows a central administrator to assign permissions to a user. To allow for the dynamic enrolment of strangers and unknown entities, we propose an entity recognition process, which is compared to traditional authentication in Figure 3. Once an entity has been recognised, i.e. has been through one cycle of the entire recognition process, trust-based security mechanisms can start up. In other words, the trust-based admission control and risk assessment components described below can operate on an entity once it has been identified through the entity recognition scheme, paralleling the way in which humans assess trust and risk once human recognition has taken place.

Recognition is therefore a necessary and sufficient requirement in the formation of trust and assessment of risk. Entity recognition is based on records about previous interactions. The more relevant the information recorded about these interactions, the more accurate the trust formation and risk assessment performed upon interactions can be.

3.3 Trust-Based Admission Control

As mentioned above, humans use the notion of trust to determine how to interact using incomplete information with known, partially-known, and unknown people. In Gray et al [3], we implemented McKnight and Chervany's trust framework and showed that a computational trust framework could be implemented such that entities might simulate human trust-based interactions by forming trust measurements and using these measurements for secure admission control. In this work, trust formation is based on the results of previous interactions with other entities. The entity can then use high level policies to specify the permitted level of admission to resources based on trust values. The results of these trials show that the trust-based admission control system reacts correctly to changes in an entity's context-specific behaviour, i.e. adjusts trust value and implements admission policies in a given context correctly according to the framework, which parallels the human trust framework.

Within this framework, when a pair of entities, p_0 and p_m, interacts for the first time, trust values have not yet been formed. To allow interaction to occur, p_0 assigns very low-risk trust-based admission rights to p_m, and from this point trust can evolve based on interactions. We foresee that a recommendation component would make this process better informed and more efficient, whereby mutual acquaintances can make recommendations to assist p_0 form an initial trust value for p_m.

Based on small world research, we surmise that in a very large mobile ad hoc network, trust can be formed and propagated between a pair of unknown entities in a predictable way through relatively short paths of mutual acquaintances. In Equation 1, we present a small world trust formula to illustrate that p_0 can indeed determine how much to trust p_m upon their first meeting. This formula forms the basis for the design of a small worlds-based recommendation component, whereby trust value certificates (TVCs) can be passed along multiple connections between entities such that initial trust values may be calculated. Because in a given context, each entity calculates trust based on the same situation-specific criteria, the trust value certificates passed between entities will be meaningful and usable.

$$Tp_0(p_m) = \frac{\sum_{k=1}^{m} w_k(Tp_{k-1}(p_k))}{m} \tag{1}$$

Where $Tp_0(p_m)$ = trust value p_0 forms for any p_m
p_0 = principal making admission control decision
p_m = principal m steps away from p_0 and requesting admission
$1 \leq 0 \leq m$ is a set of steps between connected principals
m = total number of steps connecting p_0 and p_n
k = current step
w_k = discounting factor (as k increases, w_k decreases)

To apply the formula above to, for instance, a database of film actors [21], we show the trust value one actor, Kevin Bacon, could form a trust value for

another actor, Charlie Chaplin, with whom he has never interacted. (In this context, interaction means acting in the same film together.) However, small world theory shows that these two actors can be linked through mutually shared connections to other actors. Along these connections, TVCs can be passed, as follows:

1. Charlie Chaplin was in Brother Can You Spare a Dime (1975) with Orson Welles. Therefore Orson Welles can form a trust value for Charlie Chaplin and pass it as a TVC to Colleen Camp, in step 2.
2. Orson Welles was in Hearts of Darkness: A Filmmaker's Apocalypse (1991) with Colleen Camp, so Ms. Camp can form a trust value for Orson Welles and add it to the TVC Mr. Welles passed to her.
3. Colleen Camp was in Trapped (2002) with Kevin Bacon, so Mr. Bacon has a trust value for Ms. Camp based on his own interactions with her.

Kevin Bacon can then evaluate the TVC for each step separating himself from Charlie Chaplin. Three connections occur here, and according to our formula, p_0 (Kevin Bacon) would form an initial trust value for p_3 (Charlie Chaplin) by taking the average of the sum of partial trust values based on the interactions between each pair of entities in the chain of connections between p_0 and p_3, as illustrated in Equation 2.

$$Tp_0(p_3) = \frac{(Tp_0(p_1))w_1 + (Tp_1(p_2))w_2 + (Tp_2(p_3))w_3}{3} \tag{2}$$

Each partial trust value in the sum is discounted according to how many steps it is away from p_0. This is expressed in the form of the discounting factor, w_k, which decreases as the number of steps separating entities increases.

Once $Tp_0(p_m)$ is calculated, p_0 can determine whether or not the value meets his criteria for admission. Should p_0 allow p_m admission, he may proceed with trust formation based on his own interactions with p_m.

This design addresses the difficulty in trust formation upon initial meeting between entities who are unknown to one another. At the same time, it also raises a further issue, concerning conflicting references should there be more than one trusted path connecting entities across the network. In this case, we foresee p_0 taking the most trusted of the available paths, so as to arrive at the most legitimate trust value for p_m. However, $Tp_0(p_m)$ does not distinguish interaction-based trust values from trust in an entity as a referee or recommender. Therefore, it may be necessary in future work to define a reliability factor whereby trust in correctness of recommendations is separate from other interaction-based trust calculations. The reliability factor could then be included in the TVC.

3.4 Risk Assessment

Risk is unavoidable and present in virtually every human interaction where there is uncertainty of outcomes. In many scenarios, risk can be mitigated through the use of records of every possible pattern or outcome. In this way, risk assessment

can be as precise as necessary because the maximum amount of information is available for assessing new situations based on previous patterns.

In a mobile ad hoc network, an entity regularly comes into contact with other entities with which it has never interacted. In this scenario, an entity, p_0, has no firsthand information with which to assess the risk of interacting with unknown entity, p_m.

However, as in the design of the trust-based admission control component presented in Section 3.3, we are able to design a risk assessment component based on small world theory. Assuming that risk assessment information can be passed across ad hoc networks via short paths of mutual acquaintances, p_0 may be able to assess the risk, where risk is defined as the probability of an unwanted outcome from interacting with p_m, based on recommended risk information.

Calculating the risk value for p_0's initial interaction with p_m is similar to calculating initial trust values in Section 3.3. The risk of the known parts can be accessed through the chain of connections and then used to formulate an overall small world risk assessment, as described in Equation 3.

$$Rp_0(p_m) = 1 - ((1 - Rp_0(p_1))(1 - Rp_1(p_2))(1 - Rp_2(p_3)) \dots (1 - Rp_n(p_m)) \quad (3)$$

Where $Rp_0(p_m)$ = risk assessment p_0 forms for interaction with any p_m
p_0 = principal making risk assessment
p_m = principal m steps away from p_0 and requesting interaction
$1 \leq 0 \leq m$ is a set of steps between connected principals
m = total number of steps connecting p_0 and p_n

We can apply the risk formula to the same chain of film actors in the example in Section 3.3, such that a risk value can be generated by Kevin Bacon to assess the risk of interacting with Charlie Chaplin. Along the connections between the actors, partial risk values can be assessed. Kevin Bacon can then evaluate the overall risk with interaction-based risk values for each step separating himself from Charlie Chaplin. Three connections occur here, and according to the risk formula, p_0 (Kevin Bacon) would form an initial risk assessment for p_3 (Charlie Chaplin) as follows:

1. Take the product of the complements of each of the partial risk assessments, which are based on the interactions between each pair of entities in the chain of connections between p_0 and p_3.
2. Take the complement of the resulting product to provide the final overall assessment of the risk of p_0's interaction with p_3.

According to our risk assessment design, as m increases, the level of risk also increases. Similarly, as m decreases, the lower the risk p_0 forms regarding interaction with p_m.

This design addresses the difficulty in risk assessment upon an initial meeting between entities who are unknown to one another. Upon initial interaction, p_0 can use partial risk assessments passed through a chain of mutual acquaintances between p_0 and p_m, such that initial assessment of risk of interaction may

developed with more complete information than if the partial risk assessments were not available. Similar to the trust-based admission control concern above, though, an issue arises when there is more than one path connecting two entities. In this case, we foresee p_0 taking the path with the lowest overall risk value, so as to be as cautious as possible in his decision-making. There may be scenarios, however, in which this level of caution is not desirable, e.g. where higher risk is offset by higher benefits. Therefore, in future work, we envisage the risk assessment component interfacing closely with the trust-based admission control component, whereby p_0 is permitted to make context-based choices in this regard. Moreover, linking the risk assessment and trust-based admission control components would enable the ability to assess the trust in the correctness of recommendations of risk assessment.

3.5 Trust Management

The vast number of entities with potentially different distinguishing characteristics expected to be interacting in mobile ad hoc networks and the high connectivity of entities in a small world network lead to the question of scalability of the entity recognition scheme discussed above. Large amounts of data may have to be stored, such as recognition information, information associated with trust, recommendations, observations, etc, on what may be resource-constrained devices with limited available memory. Consequently, the size of the cache, the place where recognition and trust information is stored, may be bound. For example, access to online file servers may not be provided in mobile ad hoc networks, which means that each entity has to carry any information that might be needed for secure decision-making.

To cope with scalability, we propose to 'forget' about some entities (that have been previously recognised) according to an algorithm, such that only the least critical entities are 'forgotten.' Because mobile ad hoc networks exhibit small world tendencies, we are able to design the algorithm based on small world characteristics, i.e. shortcuts and clusters. The Small-wOrld-based Forgetting Algorithm (SOFA) we propose in this regard can be helpful in the maintenance of trust-based information, depending on which source of trust is considered to be most important in the given scenario, such as recommendations and observations.

Where *recommendation data* is more important than other sources of trust, it is essential that SOFA is designed to remember entities that are 'pivots,' [22] i.e. those that have significant long-range shortcuts which span communities. In this way, the stored data will be that which is most valuable, i.e. trust information about many entities throughout the network. An algorithm such as the Shortcut Hunter mentioned above can then be used to retrieve certificate chains via these entities.

Where *observation data* is more important than other sources of trust, it is important that the algorithm is designed to remember entities with whom collaboration may occur based on the next contextual cluster. Two points are key in this regard. First, as shown in Section 3.3, it is important to be able to retrieve

trust-based information based on previous observations of an entity's behaviour. Second, in a small world, clusters of entities form according to different criteria, e.g. geographical location. For example, assuming that an entity, p_0, is roaming to another environment and knows specific information about this environment. p_0 will most likely wish to have available trust-based information about any p_m most likely to be present in the destination environment. Matsuo shows [23] that 'a cluster often shows the particular context,' and describes a Small-World Clustering algorithm to identify a cluster's context. Therefore, knowing in advance in which context p_0 is likely to be, p_0's cache should contain information relevant to p_m in the particular contextual cluster. The cluster may also be used to establish the probability of likely future collaborations [24], based on the cliquishness, or number of mutual acquaintances, within that cluster. Care should be taken regarding pivot entities, as they may not necessarily be directly related to the cluster to which they have shortcut connections.

It is important to note, that this algorithm should interface with the trust-based admission control and risk assessment components. Even if an entity has many shortcuts, it may not be trustworthy or worth the risk of interaction, and in these cases, the entity should be forgotten. Moreover, it may be useful to remember and avoid 'bad' entities, i.e. those which behave incorrectly. This raises an interesting paradox, however, in that if p_0 remembers a bad p_m, p_m may simply establish a new identity and his old identity could be forgotten, but if p_0 completely forgets about a bad p_m, p_m may retain his identity and it would have been worth remembering him. We have a possible solution to this paradox, but it is outside the scope of this paper.

3.6 Summary

In this section we described our trust-based security architecture and its associated components. We showed how existing small world principles, such as shortcuts, clustering, and distributed search algorithms, apply to the domain of ad hoc computing, specifically within our trust-based security framework. First, it allows quick trust formation and risk assessment through short chains of mutually-known entities. Second, it directs retention of information about entities in the cache, through the SOFA algorithm, thereby reducing the overall size of the cache.

4 Conclusions and Future Work

Because traditional security measures do not scale well in the envisaged massive, networked infrastructure of diverse entities partaking in collaborative applications with each other, we proposed the provision of trust-based security measures that make use of small world concepts. In this regard, we provided an overview of small world research, in which we highlighted areas that are relevant to the design of a trust-based security system. We found that self-organising networks, such as certain types of mobile ad hoc networks, exhibit small world tendencies.

Based on this premise, we are able to incorporate small world characteristics in the design of three security components for self-organising networks, based on trust, risk, and entity recognition. We then presented the design of these three components.

Our component designs are based on the concepts of recognition, trust, and risk, and each component parallels recognition, trust, and risk in human ad hoc collaborative environments. The component designs address the difficulty in entity recognition, admission control, risk assessment, and trust management upon initial meeting between entities who are unknown to one another in large, self-organising networks.

Applying existing ideas from small world research to a trust-based security architecture gives us the following results. First, it indicates that previously-unknown entities should be able to quickly establish initial trust in one another, based on short chains of recommendations via mutually-known entities. Second, it directs trust management, particularly by assisting an entity in determining which information is important to retain and which entities can be "forgotten," as demonstrated by SOFA. Next, with regard to small world influence in trust formation and risk assessment, we found that a given entity may be potentially provided with more complete information, via mutually-trusted entities, to be assessed than would be available in completely random networks (where decisions have to be made based on either direct observation or prohibitively-large recommendation chains). Having more complete information at its disposal, then, enables an entity to make more informed and predictable decisions regarding interaction with unknown entities. Increasing the informedness and predictability of decision-making enables the entire system to be more secure.

We identified future work in four key areas. First, an issue arises when there are more than one equally-short paths connecting two entities. In this scenario, we must refine the design of the trust and risk components such that an entity can evaluate equally-short paths and choose which path is most suitable. Here, we foresee the integration of criteria assessment for determining context-sensitive path suitability. Second, we determined that there is a need to distinguish interaction-based trust values from trust in an entity as a referee or recommender. Therefore, in future work, we foresee the definition of a reliability factor whereby trust in correctness of recommendations is separate from other interaction-based trust calculations. Third, we identified the need for all three designs to interface with each other, enabling them to work together in providing entity recognition, risk assessment, and trust-based admission control. In this way, relevant interaction-based and recommended information can be shared amongst each of the three components. Finally, we wish to explore possible solutions to the paradox regarding the retaining of information about 'bad' entities.

Acknowledgements

The authors would like to thank Raymond Cunningham for his valued input. This work is funded by the SECURE project (IST-2001-32486), a part of the EU FET Global Computing initiative.

References

1. Marsh, S.: Formalising Trust as a Computational Concept. PhD thesis, University of Stirling, Department of Computer Science and Mathematics (1994)
2. McKnight, D., Chervany, N.: The Meanings of Trust. MISRC 96-04, University of Minnesota, Management Informations Systems Research Center, University of Minnesota (1996)
3. Gray, E., O'Connell, P., Jensen, C., Weber, S., Seigneur, J.M., Yong, C.: Towards a Framework for Assessing Trust-Based Admission Control in Collaborative Ad Hoc Applications. Technical Report 66, Department of Computer Science, Trinity College Dublin (2002)
4. Matthews, R.: Six Degrees of Separation. World Link (2000)
5. Milgram, S.: The Small World Problem. Psychology Today **61** (1967)
6. Travers, J., Milgram, S.: An Experimental Study of the Small World Problem. Sociometry **32** (1969) 425–443
7. Watts, D., Strogatz, S.: Collective Dynamics of 'Small-World' Networks. Nature **393** (1998) 440–442
8. Watts, D.: Small Worlds, The Dynamics of Networks Between Order and Randomness. Princeton University Press (1999)
9. Watts, D., Strogatz, S.: Kevin Bacon, the Small-World, and Why It All Matters. http://www.santafe.edu/sfi/publications/Bulletins/bulletinFall99/workInProgress/smallWorld.html (1999)
10. Reynolds, P.: Oracle of Bacon. (http://www.cs.virginia.edu/oracle/)
11. Adamic, L.: The Small World Web. In Abiteboul, S., Vercoustre, A.M., eds.: Proc. 3rd European Conf. Research and Advanced Technology for Digital Libraries, ECDL. Number 1696, Springer-Verlag (1999) 443–452
12. Kleinberg, J.: The Small-World Phenomenon: An Algorithmic Perspective. In: Proc. of the 32nd ACM Symposium on Theory of Computing. (2000)
13. Watts, D., Dodds, P., Newman, M.: Identity and Search in Social Networks. Science **296** (2002) 1302–1305
14. Hubaux, J.P., Buttyan, L., Capkun, S.: The Quest for Security in Mobile Ad Hoc Networks. In: Proc. of the ACM Symposium on Mobile Ad Hoc Networking and Computing (MobiHOC). (2001)
15. Capkun, S., Buttyan, L., Hubaux, J.P.: Small Worlds in Security Systems: an Analysis of the PGP Certificate Graph. In: New Security Paradigms Workshop, Norfolk, VA (2002)
16. Stajano, F.: Security for Ubiquitous Computing. Wiley (2002)
17. Seigneur, J.M., Farrell, S., Jensen, C.D.: Secure Ubiquitous Computing Based on Entity Recognition. In: Ubicomp'02 Security Workshop, Gothenburg (2002)
18. Seigneur, J.M., Farrell, S., Jensen, C.D., Gray, E., Yong, C.: End-to-end trust in pervasive computing starts with recognition. In: Proceedings of the First International Conference on Security in Pervasive Computing, Boppard, Germany (2003 [to appear])

19. ITU: Information Technology - Opens Systems Interconnection - The Directory: Authentication Framework. Number X.509 in ITU–T Recomandation. International Telecomunication Union (1993)
20. Kohl, J., Neuman, B.: The Kerberos Network Authentication Service (Version 5). RFC 1510, IETF (1993)
21. IMDB: Internet Movie Database. (http://www.imdb.com)
22. Venkatraman, M., Yu, B., Singh, M.: Trust and Reputation Management in a Small-World Network. Technical report (2002)
23. Matsuo, Y.: Clustering Using Small World Structure. In: Knowledge-Based Intelligent Information and Engineering Systems, Crema, Italy (2002)
24. Newman, M.: Clustering and Preferential Attachment in Growing Networks. Phys. Rev. E **64** (2001)