On Trust Establishment in Mobile *Ad-Hoc* Networks

Laurent Eschenauer, Virgil D. Gligor, and John Baras*

Electrical and Computer Engineering Department, University of Maryland
College Park, MD 20742, USA
{laurent, gligor, baras}@eng.umd.edu

Abstract. We present some properties of trust establishment in mobile, ad-hoc networks and illustrate how they differ from those of trust establishment in the Internet. We motivate these differences by providing an example of ad-hoc network use in battlefield scenarios, yet equally practical examples can be found in non-military environments. We argue that peer-to-peer networks are especially suitable to solve the problems of generation, distribution, and discovery of trust evidence in mobile ad-hoc networks, and illustrate the importance of evaluation metrics in trust establishment.

1 Introduction

We view the notion of "trust" among entities engaged in various protocols as a set of relations established on the basis of a body of supporting assurance (trust) evidence. In traditional networks, most trust evidence is generated via potentially lengthy assurance processes, distributed off-line, and assumed to be valid on long terms and certain at the time when trust relations derived from it are exercised. Trust relations established as a consequence of supporting trust evidence are often cached as certificates and as trust links (e.g., hierarchical or peer links) among the principals included in these relations or among their "home domains." Both certificates and trust relations are later used in authorizing client access to servers.

In contrast, few of these characteristics of trust relations and trust evidence are prevalent in *mobile ad-hoc networks (MANETs)*. Lack of a fixed networking infrastructure, high mobility of the nodes, limited-range and unreliability of wireless links are some of the characteristics of MANET environments that constrain the design of a trust establishment scheme. In particular, trust relations may have to be established using only on-line-available evidence, may be short-term and largely peer-to-peer, where the peers may not necessarily have a

* This work was supported in part by U.S. Army Research Office under Award No. DAAD19-01-1-0494, and by the U.S. Army Research Laboratory under Cooperative Agreement DAAD19-01-2-0011 for the Collaborative Technology Alliance for Communications and Networks.

B. Christianson et al. (Eds.): Security Protocols 2002, LNCS 2845, pp. 47–66, 2004.

relevant "home domain" that can be placed into a recognizable trust hierarchy, and may be uncertain.

In this paper we argue that in MANETs a substantial body of trust evidence needs to be (1) generated, stored, and protected across network nodes, (2) routed dynamically where most needed, and (3) evaluated "on the fly" to substantiate dynamically formed trust relations. In particular, the management of trust evidence should allow alternate paths of trust relations to be formed and discovered using limited backtracking though the ad-hoc network, and should balance between the reinforcement of evidence that leads to "high-certainty" trust paths and the ability to discover alternate paths. We derive several design parameters for the generation and distribution of trust evidence in MANETs by analyzing salient characteristics of peer-to-peer file sharing protocols.

2 On Trust Establishment Differences between the Internet and MANETs

In this section, we review some of the basic notions of trust establishment and explore how these notions differ in the MANET environment from those in the Internet environmnet. We also derive a set of requirements for trust establishment in MANETs. Much of the theory underlying the presentation of basic notions can be found in Maurer [16], Kohlas and Maurer [13], others [15] [9]. We focus exclusively on some empirical properties of evidence for trust establishment that help differentiate the traditional Internet notions from those of MANETs.

2.1 Basic Notions of Trust Establishment

We view the process of trust establishment as the application of an evaluation metric to a body of trust evidence. The outcome of the trust establishment process is a trust relation. The evidence may be obtained on- or off-line and may include already established trust relations. An established trust relation constitutes evidence that can be used in other trust establishment processes, and can be composed with other relations to form more abstract or more general trust relations. The composition of trust relations usually requires the composition of evidence and of evidence evaluations.

An Example of Authentication-Trust Establishment. Consider the trust relation[1] "A accepts B's authentication of X", which is established between principals A, B, and X. This relation is established as the composition of two basic relations resulting from two separate trust-establishment processes; i.e., "certification authority B accepts X's authentication evidence," and "certification authority A accepts B's authentication of any principal registered by B". The first relation may be established by principal B's off-line evaluation of a body of trust evidence presented by principal X. For example, B may require several

[1] Although we focus on authentication, similar notions can be defined for trust establishment in the access control arera.

pieces of evidence attesting to X's identity. Specifically, B may require two pieces of authentication evidence from the following set: driver license, passport, employment identity card, documentation indicating current property ownership or credit-line activity. Once the trust relation is established, it is cached as (1) a certificate signed by B associating X's public key with X, and (2) a relation stored in B's "trust database" registering principal X with B. The domain of certification authority B becomes X's "home domain."

The second relation, namely "certification authority A accepts B's authentication of any principal registered by B," may be established by principal A's *off-line* evaluation of a body of trust evidence presented by principal B indicating that:

- certification authority B's authentication of the principals registered with it (e.g., X) is done using "acceptable" mechanisms and policies; and
- certification authority B's registration database, which includes principal X's registration, is protected using "acceptable" mechanisms and policies;
- certification authority B's server is managed using "acceptable" administrative, physical, and personnel policies;
- certification authority B does not have skills and interests that diverge from those of A.

Evidence regarding the "acceptability" of various mechanisms and policies is collected off-line, using potentially lengthy assurance procedures, such as those prescribed by the Common Criteria's assurance evaluation levels [8]. Certification authority A uses an evaluation metric to determine whether B's authentication mechanisms and policies are (at least) as good as his own, and the evidence used by the metric is *stable* and *long-term*. Evidence is stable if the authentication mechanisms and policies used by B do not change, either intentionally or accidentally, unbeknownst to A. Evidence is long-term, if it lasts at least as long as the process of gathering and evaluating assurance evidence, which can be of the order of weeks or months. After the trust relation "certification authority A accepts B's authentication of any principal registered by B" is established by A, it is cached (1) as a certificate associating B's public key with B that is signed by A, and (2) as a relation stored in A's "trust database" registering principal B with A. The domain of certification authority A becomes B's "home domain."

Transitivity of Trust Establishment. Trust relation "certification authority A accepts B's authentication of any principal registered by B" is clearly *reflexive* since A accepts its own authentication of principals it registers. However, should it be *transitive*? That is, should the trust establishment process be transitive? For example, if "A accepts B's authentication of any principal registered by B" and "B accepts Y's authentication of principal Z registered by Y," does it mean that "A accepts Y's authentication of principal Z registered by Y"? And if so, does this hold for any principals Y and Z?

Before accepting that transitivity should hold, A uses his "evaluation metric" to determine two properties of evidence. First, A determines that B's evaluation of Y's body of evidence is the same as (or stronger than) A's evaluation of B's

body of evidence (viz., example above). Second, A determines that B's trust relation with Y is (at least) as stable and long-term as his A's own with B. If these two properties of evidence hold for all Y's and Z's, then the more general trust relation "A accepts Y's authentication of any principal" should also hold. In practice, this general trust relation would hold for all Y's whose home domains are sub-domains of B's home domain. This is the case because B would control the adequacy, stability, and duration of Y's authentication mechanisms and policies, and hence could provide the evidence that would satisfy A's evaluation metric. However, evidence regarding Y's authentication mechanisms and policies may not pass A's evaluation metric, and A would not accept Y's authentication of any principal. For example, the evidence used in establishing B's trust relation with Y may be short-lived or unstable. In this case, Y could change its authentication policies, thereby invalidating evaluated evidence, unbeknownst to A and B. A would want to be protected from such events by denying transitivity regardless of whether B accepts Y's authentication of Z.

The principal characteristics of evidence used to establish transitive trust in the example given above are "uniformity" and "availability." Uniformity means that all evidence used to establish transitive trust satisfied the same, global, "metrics" of adequacy, stability, and long-term endurance. Availability means that all evidence could be evaluated either on-line or off-line at any time by a principal wishing to establish a trust relation.

Uncertainty in Trust Establishment. Transitive trust formed the basis for the definition of simple trust hierarchies, possibly interconnected by "peer" links. All early system designs supporting such hierarchies assumed either implicitly [15] or explicitly [9] that evidence for recommending trust from principal to principal was "uniform" and "available." In contrast, starting with Yahalom *et al.* [24], it was realized that, in general, trust evidence need not be uniform and hence could be uncertain. Pretty Good Privacy (PGP) [25] provides the first practical example where some "uncertainty" is allowed in authentication, although PGP does not support transitive trust. Later work by Kohlas and Maurer [13] formalizes the notion of evidence uncertainty and provides precise and fairly general principles for evaluating trust evidence.

Guaranteed Connectivity to Trust-Infrastructure Servers. To be scalable, Public Key Infrastructures (PKIs) establish trust among certification authorities rather than among individual principals. Transitive trust relations among certification authorities allows us to establish authentication trust among principals registered by different certification authorities, since it allows the traversal of certification authorities separating pairs of principals; i.e., the traversal of trust paths. Traversal of trust paths does not require that certification authorities be on-line permanently. Certification authorities store certificates in directories associated with "home domains" whenever trust relations are established, and hence directory hierarchies mirror trust hierarchies. Therefore, directory servers must be available and on-line permanently to enable trust path traversals by any principal at any time, whereas certification authority servers need be on-line only when trust relations are established and certificates are signed and stored in di-

rectories. Nevertheless, principals establishing trust relations or traversing directory hierarchies to establish, or verify the validity of, trust paths need guaranteed communication connectivity to certification authority and directory servers.

2.2 Internet *vs.* Mobile *Ad-Hoc* Networks

Ad-hoc networking refers to spontaneous formation of a network of nodes without the help of any infrastructure, usually through wireless communication channels. In *ad-hoc* networks, a basic routing infrastructure emerges through the collaboration of every node with its neighbors to forward packets towards chosen destinations. This basic infrastructure is highly dynamic not just because of node mobility (which also characterizes mobile IP networks) but also because of lack of guaranteed node connectivity (which is not necessarily a characteristic of mobile IP networks). In *ad-hoc* networks, lack of guaranteed connectivity is caused by the limited-range, potentially unreliable, wireless communication. The absence of a routing infrastructure that would assure connectivity of both fixed and mobile nodes precludes supporting a stable, long-term, trust infrastructure, such as a hierarchy of trust relations among subsets of network nodes. It also constrains the trust establishment process to short, fast, on-line-only protocols using only subsets of the established trust relations, since not all nodes that established trust relations may be reachable.

Trust Establishment without a Trust Infrastructure. In general, the Internet relies on a fixed trust infrastructure of certification-authority and directory servers for both fixed and mobile nodes (i.e., Mobile IPv6 nodes). These servers must be available on-line and reachable by principals when needed; e.g., certification authority servers, when certificates are created and signed, and directory servers permanently.

In contrast, a fixed infrastructure of certification-authority and directory servers may not always be reachable in a MANET (viz. Section 3, scenarios 2 and 3). This is because MANETs cannot assure the connectivity required to these servers; e.g., both a mobile node and the foreign-domain nodes with which it communicates can be disconnected from the directory server storing the certificates defined in that node's home domain [2]. Therefore, MANETs cannot rely exclusively on trust relations that are represented as certificates stored in directory hierarchies, since connectivity to the required servers may not be available when needed. MANETs must support *peer-to-peer relations* defined as the outcomes of any principal's evaluation of trust evidence from *any* principals in the network, and must store these trust relations in the nodes of the *ad-hoc* network.

Short-lived, Fast, and On-line-only Trust Establishment. In the Internet, trust relations are established for the long term and are stable. This is possible if security policies and assurances do not change very often and therefore do not need to be re-evaluated frequently.

[2] Note that this is not the case for mobility in the Internet. Mobile IPv6 takes care of roaming by providing a "care of" address bound to the actual mobile address. This solution is not possible for MANETs since the home of a node and its "care of" address may be physically unreachable.

In contrast, there is little long-term stability of evidence in MANETs. The security of a mobile node may depend of its location and cannot be a priori determined. For example, node capture by an adversary becomes possible and probable in some environments such as military battlefields. Trust relations involving a captured node need to be invalidated, and new trust evidence need to be collected and evaluated to maintain node connectivity in the *ad-hoc* network. Therefore, trust relations can be short-lived and the collection and evaluation of trust evidence becomes a recurrent and relatively frequent process. This process has to be fast to avoid crippling delays in the communication system; e.g., two mobile nodes may have a short time frame to communicate because of wireless range limitations, and trust establishment should not prevent these nodes from communicating securely by imposing a slow, lengthy process. To be fast, the trust establishment process may have to be executed entirely on-line since off-line collection and evaluation of evidence is impractical; e.g., visually verifying an identity document is not possible.

Trust Establishment with Incomplete Evidence. In the Internet, it is highly improbable that some trust relation remains unavailable for extended periods of time (e.g., a certificate verification on a trust path cannot performed for a day) due to connectivity failures. Network connectivity is guaranteed through redundancy of communication links, and routes and servers are replicated to guarantee availability. In general, it is fair to assume that the entire body of evidence necessary for trust establishment is available in the Internet when needed. In contrast, node connectivity is not guaranteed in MANETs and all established evidence cannot be assumed to be available for all nodes all the time. Trust establishment has to be performed with incomplete and hence uncertain trust evidence.

In summary, trust establishment in MANETs requires protocols that are:

- peer-to-peer, independent of a pre-established trust infrastructure (i.e., certification authority and directory servers);
- short, fast, and on-line; and
- flexible and support uncertain and incomplete trust evidence.

3 An Example with Three Scenarios

In this section we present an example to motivate the requirements of trust establishment presented above. The example consists of three related battlefield scenarios. For the sake of brevity, we omit relevant examples from non-military applications.

3.1 Scenario 1

In Figure 1 we illustrate a battlefield environment in which units of coalition of United States (US) and United Kingdom (UK) forces performs separate operations. To support these operations, various communication systems are involved,

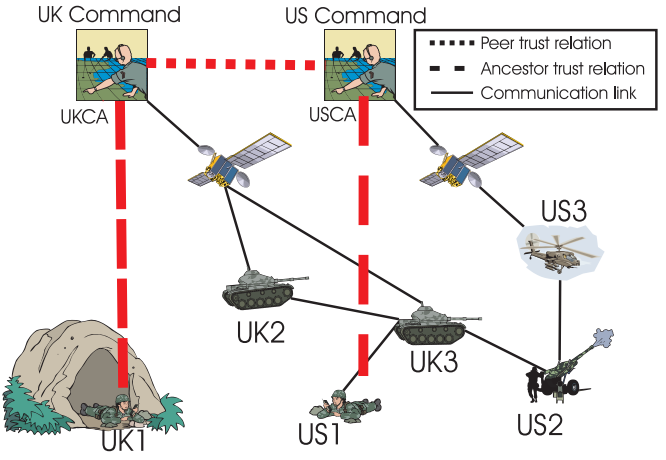

Fig. 1. A battlefield scenario. UK1 is lost and can only communicate with US1

ranging from short-range wireless (e.g., for infantry), to long-range directional wireless links (e.g., used between artillery pieces), and to satellite communication (e.g., connecting the battlefield with the US and UK operation commands). In this scenario, assume that a British unit (UK1) is lost and takes refuge in a nearby cave. UK1 needs to call for backup, but the only unit in communication range is an American unit (US1) taking part in a different operation than that of UK1. The British unit, UK1, has to authenticate itself to US1 to get access to the *ad-hoc* US network and call the UK operations command for help. UK1 requests access to the *ad-hoc* US network and presents an identity certificate signed by UKCA, the British certification authority. The US network access policy requires that any accessor presents a valid identity certificate from a US-recognized and trusted authority. Node US1 needs to decide whether the node claiming to be UK1 should be allowed access to the *ad-hoc* US network. To decide wether UK1's certificate is valid, US1 contacts the directory server at US operations command and obtains a UKCA certificate signed by USCA, the US certification authority. US1 and accepts USCA's signature on the UKCA's certificate, then accepts UKCA's signature on UK1's certificate, thereby exercising the transitive trust relations established between the US and UK operations commands and their respective units. Node US1 grants access to the *ad-hoc* US network to UK1. Note that the established trust infrastructure of the Internet helps solve UK1's problem, since all necessary trust relations (i.e., evaluated evidence) are available on-line.

3.2 Scenario 2

Assume that, due to inclement weather conditions, satellite links are unavailable. When US1 receives UK1's request and certificate signed by UKCA, it can't

contact its operations command center to retrieve UKCA's certificate from a directory server, and therefore it cannot verify the signature on UK1's certificate. However, suppose that a couple hours ago while in a different operation, a US helicopter unit, US3, visually identified the lost British unit, UK1. US3 could have *proactively* generated a certificate for UK1 and made it available in the *ad-hoc* US network. Alternately, US3 could generate and sign a certificate for UK1 now. This certificate is the only piece of evidence that could allow authentication of UK1 by US1. However, currently there is currently no scheme to specify how and when such a certificate is generated, how it can be distributed to others nodes in the network, how it should be evaluated by US1 to take its access decision and, finally, how it can be revoked by US3, if the need arise. In section 4 we present our approach on how to solve these issues.

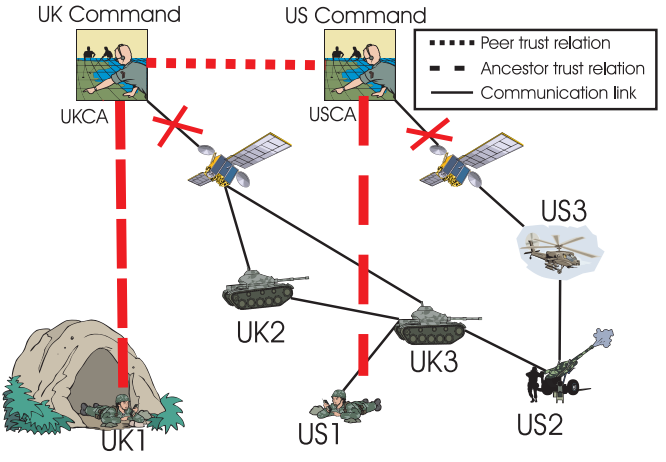

Fig. 2. A battlefield scenario. UK1 is lost and can only communicate with US1. The satellite links are down due to inclement weather

3.3 Scenario 3

Figure 3 illustrates a United Nations humanitarian convoy (UN1) that is approaching and preparing to cross a bridge separating two battlefield "zones". Before crossing the bridge to enter the new zone, UN1 must request a "zone report" from nearby military units to verify that the zone is safe. UN1 sends a request for a zone report and attaches its credentials (Table 1.b) as authentication evidence to the request. A British unit, UK3, receives the request and is in a position to issue a zone report. However, to issue the zone report, UK3 needs to apply its evaluation metric (Table 1.d and 1.e) to the presented evidence (and

the evidence already in its possession by other means) and to verify that it satisfies the policy it must enforce for providing zone reports (Table 1.a). However,

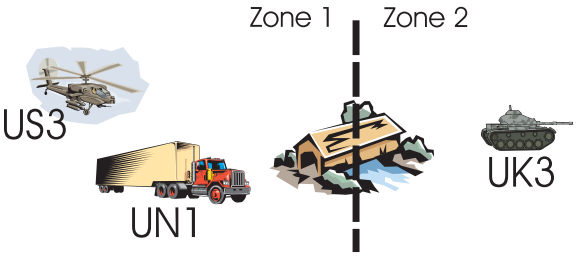

Fig. 3. A battlefield scenario

UK3 has a limited set of already established trust relations (Table 1.c) and it is not hard to see that some evidence provided by UN1 (1) is useful but cannot be verified (i.e., certificates signed by USCA and US3 cannot be verified by UK3 since it does not have a direct trust relation to USCA and US3 and the satellite links are unavailable); or (2) can be verified but is not useful (i.e., GPS1 is trusted to provide location information but the UK3 evaluation metric rates any GPS source to provide only low-confidence information whereas high-confidence information is required by the UK3 policy). Therefore, UK3 needs to collect and evaluate evidence regarding USCA and US3 using the *ad-hoc* network only, since the central directory at its operation command remains unavailable.

4 Our Approach

4.1 Generation of Trust Evidence

In our approach, any node can generate trust evidence about any other node. Evidence may be an identity, a public key, a location, an independent security assessment, or any other information required by the policy and the evaluation metric used to establish trust. Evidence is usually obtained off-line (e.g. visual identification, audio exchange [2], physical contact [20] [21], etc.), but can also be obtained on-line. When a principal generates a piece of evidence, he signs it with its own private key, specify its lifetime and makes it available to other through the network. PGP is an instance of this framework, where evidence is only a public key.

A principal may revoke a piece of evidence it produced by generating a revocation certificate for that piece of evidence and making it available to others, at any time before the evidence expires. Moreover, a principal can revoke evidence generated by others by creating contradictory evidence and distributing it. Evidence that invalidates other existant evidence can be accumulated from multiple,

Table 1. An Example of a Policy Statement, Evaluation Metric, and Credentials and Trust Relations

a. UK3's policy for providing "zone reports":
(*Role* = UK/US military ∨ UN convoy) with confidence = medium
∧(*Location*=neighbors) with confidence = high
b. UN1's request presents credentials:
Cert(*Role*=UNConvoy)$_{USCA}$
Cert(*Location*/GPS=zone2)$_{GPS1}$
Cert(*Location*/Visual=zone2)$_{US3}$
c. UK3's trust relations:
UKCA for *Role*; GPS1, UAV1, and UK1 for *Location*
d. UK3's metric for confidence evaluation of location evidence
Type(source) = GPS and source trusted → confidence = low
Type(source) = UAV and source trusted → confidence = low
Type(src1) = UAV ∧ Type(src2) = GPS
and src1 and src2 trusted → confidence = medium
Type(source) = Visual and source trusted → confidence = high
Other → confidence = null
e. UK3's metric for confidence evaluation of role evidence:
Type(source) = CA and source trusted → confidence = high
Other → confidence = null

independent, and diverses sources and will cause trust metrics to produce low confidence parameters.

It may seem dangerous to allow anyone to publish evidence within the *ad-hoc* network without control of any kind. For example, a malicious node may introduce and sign false evidence thereby casting doubt about the current trust relations of nodes and forcing them to try to verify the veracity of the (false) evidence. To protect against malicious nodes, whenever the possibility of invalidation of extant trust evidence (e.g., evidence revocation) arises, the policy must require redundant, independent pieces of (revocation) evidence from diverse sources before starting the evaluation process. Alternatively, the evaluation metric of the policy may rate the evidence provided by certain nodes as being low-confidence information. In any case, the policy and its evaluation metric can also be designed to protect against false evidence.

4.2 Distribution of Trust Evidence

Characteristics. Every principal is required to sign the pieces of evidence it produces. A principal can distribute trust evidence within the network and can even get disconnected afterwards. A producer of trust evidence does not have to be reachable at the time its evidence is being evaluated. Evidence can be replicated across various nodes to guarantee availability. This problem of evidence availability is similar to those that appear in distributed data storage systems, where

information is distributed across multiple nodes in a network, and a request for a piece of stored information is dynamically routed to the closest source.

However, trust evidence distribution is more complex than a simple "request routing" problem. A principal may need more than one answer per request, and hence *all* valid answers to a request should ideally be collected. For example, `REQUEST(Alice/location)` should return all pieces of evidence about the location of Alice. Typical distributed data storage systems do not return all valid requests; e.g. `REQUEST(my_song.mp3)` would return one file even if there are multiple versions of my_song each having different bit rates and length. Moreover a principal may simply not know what evidence to request, and hence wildcard requests have to be supported; e.g. `REQUEST(Alice/*)` should return all pieces of evidence about Alice available in the network.

Peer-to-peer file sharing for evidence distribution. The problem of evidence distribution shares many characteristics of distributed data storage systems, and yet is different. It is interesting to examine current peer-to-peer, file-sharing systems to understand their characteristics and limitations regarding trust evidence distribution. Peer-to-peer networking has received a lot of interest attention recently, particularly from the services industry [10] [17], the open-source [7], and research communities [1] [14] [22]. They evolved from very simple protocols, such as Napster (which uses a centralized index) and Gnutella (which uses request flooding) to more elaborate ones, such as Freenet (which guarantees request anonymity and uses hash-based request routing) [7] and Oceanstore (which routes requests using Plaxton trees) [14].

We analyzed Freenet as a tool for evidence distribution because of the characteristics of its request routing architecture. In particular, in Freenet requests are routed in the network instead of flooding. The routing is based on a hash of the requested keywords. Files are replicated by caching at every node. Frequently requested files are highly replicated across the network while file that are rarely requested are slowly evicted from caches. Anonymity and secrecy are guaranteed. It is not possible to know which node is requesting which file, and it is not easy to discover where a particular file is stored.

Request routing in Freenet is adaptive and improves with time; combined with the caching policy it shows an interesting locality property: information converges where needed and is forgotten where not requested. This suits particularly well the locality property of trust establishment in the MANET (a node tends to establish trust with nearby neighbors). This optimized routing allows faster distribution and revocation of pieces of evidence. However, the Freenet approach does not support wildcard requests and provides only one answer per request (due to the nature of its routing mechanism). Moreover, access to various sources of information evolves only by path reinforcement. As a consequence, some sources of information providing non-usable data are reinforced, and other sources are not discovered. The reinforcement strategy of Freenet does not preserve the diversity of information sources in the network. A new system has to be designed that shares the advantages of Freenet without exhibiting its drawbacks.

Swarm intelligence for trust evidence distribution. Swarm intelligence is a framework developed from the observation of ants' colonies. While a single ant is a very simple insect, groups of ants can cooperate and solve complex problems such as finding the shortest path to a food source or building complex structures. Ants do not communicate directly with each other; instead they induce cooperation by interacting with their environment (e.g., leaving a pheromone trail). When trying to find an optimum solution (e.g., shortest path to food source), cooperation leads to reinforcement of good solutions (positive feedback); more over, the natural decay of a pheromone trail enables regulation (negative feedback) that helps discover of new paths.

Numerous algorithms have been developed from these observations and applied to problems such as the traveling salesman, graph coloring, routing in networks [6],... Swarm intelligence is particularly suited for solving optimization problems in dynamically changing environments such as those of MANETs because of the balance between positive feedback that helps reinforce a good solution and the regulation process that enables discovery of new solutions appearing because of changes in the environment.

The problem of discovering proper sources of trust evidence in a MANET (and the problem of resource discovery in a network in general) is similar to the discovery of food supplies for an ant colony. It requires exploration of the environment with reinforcement of good solutions but also regulation that allows new sources to be discovered.

4.3 Application of an Evaluation Metric to a Body of Evidence

In specifying a trust management policy, we distinguish between a *policy decision* and a *trust metric* for practical rather than fundamental reasons. A metric is used to assign a confidence value to pieces of evidence of the same nature[3]. For instance, if we have three sources of evidence providing three different locations for Alice, how do we determine Alice's actual location and how confident are we of that determination? In contrast, a policy decision is a local procedure which, based on a set of evidence parameters and their required confidence value, outputs the outcome of the decision. In practice, policy decisions are locally enforced but may be based on trust metrics shared by other local policies. Similarly, the same policy decision may use different trust metrics (as in the case of UK3's metrics in Scenario 3 above) for different parameters. Different types of policy decisions have been proposed that apply a policy to a set of credentials and output a decision [4] [5].

Trust metrics to evaluate uncertain and incomplete sets of evidence has been an active field of research. Different "trust metrics" have been developed [16] [18] [24] and properties of these metrics have been studied [13]. However, the only practical trust metric developed and implemented has been the

[3] Different metrics may be used for different type of evidence (e.g. one may use a discrete level metric to characterize confidence in location, but a continuous metric to characterize confidence in a public key).

one of PGP [25]. Based on a very limited notion of uncertainty, this metric handles only the evaluation of trust in a chain of keys, with limited "levels of trust" (i.e. untrusted, marginal, full). There is a need to develop new trust metrics that apply to different types of evidence, not just chains of keys, are fine-grained in the sense that output wide set of uncertainty levels, and are flexible, in the sense that they can apply to incomplete sets of evidence.

5 Related Work

5.1 Pretty Good Privacy

In PGP [25], any user can sign another user's key. These signatures form a network of peer trust relations, often described as the *web of trust* [25]. The confidence in a trust path between two nodes of the web of trust is evaluated via a simple metric consisting of 4 "levels of trust" and a set of rules (e.g.: a key is marginally trusted if signed by two independent, marginally trusted, keys).

Although the PGP web of trust is fully peer-to-peer in its concepts, it is not in implementation. Public keys are published in *key servers* [19] maintaining a database of keys and discovering trust paths amongst them. This solution is efficient for the Internet but not possible for the MANET since there is no guaranteed connectivity with a key server. Hubaux *et al.* [12] propose a distributed implementation of PGP where each user stores a subset of the trust graph and proceeds to fusion of his set with other users' sets to discover trust path.

The trust metric implemented in PGP is simple and can lead to counter intuitive decision being made, as discussed by Maurer [13].

5.2 IBM's Trust Establishment System

IBM Research Laboratory developed a trust establishment framework [11] allowing the "bottom-up" emergence of a public-key infrastructure through exchange of certificates, containing various pieces of evidence about principals, and evaluation of these by a *Trust Policy Language*. When certificates about a principal are missing, they are automatically collected from peer servers. The policy language supports negative certificates, which allows complex non-monotonous policies. However, the trust policy language does not support *uncertain evidence* explicitly; as this is considered part of the policy specification.

This work is targeted to the Internet, where connectivity is guaranteed between servers. Missing certificates are collected from peer servers (either known *a priori* or referenced in other certificates). The collection mechanism is not suitable for the MANET environment were connectivity is not guaranteed. Our peer-to-peer evidence distribution mechanism would be a suitable solution to replace the certificate repositories and support the IBM's *trust engine* to provide a full peer-to-peer implementation.

5.3 The Resurrecting Duckling

Stajano and Anderson's resurrecting duckling [20] and its descendants [2] [21] represent a peer-to-peer trust establishment framework in which principals authenticate their communication channel by first exchanging keying material via an out-of-band physical contact. The goal of this approach is different from ours; i.e., it is not intended to provide peer-to-peer entity authentication, nor is it intended to handle uncertain evidence. The established trust is binary: the communication channel is either secure or is not.

6 Conclusions and Future Work

The notion of trust establishment in mobile *ad-hoc* networks (MANETs) can differ from that in the (mobile) Internet in fundamental ways. Specifically, it has the trust establishment process has to be (1) peer-to-peer, (2) short, fast, and on-line-only, and (3) flexible enough to allow uncertain and incomplete trust evidence.

We presentend a framework for trust establishment that supports the requirements for MANETs and relies on peer-to-peer file-sharing for evidence distribution through the network. The problem of evidence distribution for trust establishment is somewhat different than the usual file sharing problem in peer-to-peer networks. For this reason, we proposed to use a swarm intelligence approach for the design of trust evidence distribution instead of simply relying on an ordinary peer-to-peer, file-sharing system. In future work, we plan to evaluate the performance of swarm-based algorithms for trust evidence distribution and revocation in a MANET environment.

Finally, we also argued that the design of metrics for the evaluation of trust evidence is a crucial aspect of trust establishment in MANETs. In future work, we plan to develop a trust management scheme integrating the confidence valuation of trust evidence with real-time, policy-compliance checking.

Disclaimer

References

1. O. Babaoglu, H. Meling, and A. Montresor, "Anthill: A Framework for the Development of Agent-Based Peer-to-Peer System", Technical Report UBLCS-2001-09, University of Bologna, Italy.
2. D. Balfanz, D.K. Smetters, P. Stewart, and H. Chi Wong, "Talking To Strangers: Authentication in Ad-Hoc Wireless Networks", in Proc. of the ISOC 2002 Network and Distributed Systems Security Symposium, February 2002.

3. T. Beth, M. Borcherding, and B. Klein, "Valuation of trust in open networks", in Proc. of ESORICS 94. Brighton, UK, November 1994.
4. M. Blaze, J. Feigenbaum, and J. Lacy, "Decentralized trust management", in Proc. of the 1996 IEEE Symposium on Security and Privacy, pages 164–173, May 1996.
5. Matt Blaze, Joan Feigenbaum, and Angelos D. Keromytis, "KeyNote: Trust management for publickey infrastructures", in Proc. Cambridge 1998 Security Protocols International Workshop, pages 59–63, 1998.
6. E. Bonabeau, M. Dorigo and G. Theraulaz, *Swarm Intelligence: From Natural to Artificial Systems*, Santa Fe Institute on the Sciences of Complexity, Oxford University Press, July 1999.
7. I. Clarke, O. Sandberg, B. Wiley, and T. W. Hong, "Freenet: A Distributed Anonymous Information Storage and Retrieval System," in Proc. of the International Computer Science Institute (ICSI) Workshop on Design Issues in Anonymity and Unobservability, Berkeley, CA, 2000.
8. *Common Criteria for Information Technology Security Evaluation – Part 3: Security Assurance Requirements*, version 2.0, CCIB-98-028, National Institute of Standards and Technology, May 1998. *http://niap.nist.gov*
9. V. D. Gligor, S.-W. Luan, and J. N. Pato, "On inter-realm authentication in large distributed systems," in Proc. of the 1992 IEEE Symposium on Research in Security and Privacy, May 1992.
10. Gnutella, http://www.gnutellanews.com/
11. A. Herzberg, Y. Mass, J. Mihaeli, D. Naor, and Y. Ravid, "Access Control Meets Public Key Infrastructure, Or: Assigning Roles to Strangers," in Proc. of the 2000 IEEE Symposium on Security and Privacy, 14-17 May 2000, Berkeley, California, USA, pages 2-14
12. J.-P. Hubaux, L. Buttyan and S. Capkun, "The Quest for Security in Mobile Ad Hoc Networks," in Proc. of the ACM Symposium on Mobile Ad Hoc Networking and Computing (MobiHOC 2001).
13. R. Kohlas and U. Maurer, "Confidence Valuation in a Public-key Infrastructure Based on Uncertain Evidence," in Proc. of Public Key Cryptography 2000, Lecture Notes in Computer Science, vol. 1751, pp. 93-112, Jan 2000.
14. J. Kubiatowicz, D. Bindel, Y. Chen, S. Czerwinski, P. Eaton, D. Geels, R. Gummadi, S. Rhea, H. Weatherspoon, W. Weimer, C. Wells, and B. Zhao, "OceanStore: An Architecture for Global-Scale Persistent Storage," in Proc. of the Ninth international Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS 2000), November 2000.
15. B. W. Lampson, M. Abadi, M. Burrows, and Edward Wobber, "Authentication in distributed systems: Theory and practice,", ACM Transactions on Computer Systems, 10(4):265–310, November 1992.
16. U. Maurer, "Modelling a Public-Key Infrastructure." in Proc. ESORICS '96 (4th European Symposium on Research in Computer Security), Rome, LNCS 1146, Springer-Verlag, Berlin 1996, 325–350.
17. Napster, http://www.napster.com
18. M. K. Reiter and S. G. Stubblebine, "Toward acceptable metrics of authentication," in Proc. of the IEEE Conference on Security and Privacy, Oakland, CA, 1997.
19. M. K. Reiter and S. G. Stubblebine, "Path independence for authentication in large-scale systems," in Proc. of the 4th ACM Conference on Computer and Communications Security, April 1997.

20. F. Stajano and R. Anderson, "The Resurrecting Duckling: Security Issues for Adhoc Wireless Networks," in Proc. of the 8th International Workshop on Security Protocols, Lecture Notes in Computer Science, Springer-Verlag, Berlin, Germany, 1999.
21. F. Stajano, "The resurrecting duckling – What next?," in Proc. of the 8th International Workshop on Security Protocols, Lecture Notes in Computer Science, Springer-Verlag, Berlin, Germany, April 2000.
22. I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan, "Chord: A scalable peer-to-peer lookup service for Internet applications," in Proc. of the 2001 ACM SIGCOMM Conference, San Diego, CA, 2001, pages 149–160.
23. E. Wobber, M. Abadi, M. Burrows, and B. Lampson, "Authentication in the Taos operating system," ACM Transactions on Computer Systems, 12(1):3–32, Feb. 1994.
24. R. Yahalom, B. Klein, and T. Beth. "Trust relationships in secure systems—A distributed authentication perspective," in Proc. of the 1993 IEEE Symposium on Research in Security and Privacy, pages 150–164, May 1993.
25. P. R. Zimmermann, *The Official PGP User's Guide*, MIT Press, 1995. (http://www-mitpress.mit.edu/ mitp/recent-books/comp/pgp-user.html)
26. L. Zhou and Z. Haas, "Securing ad hoc networks," IEEE Network, 13(6):24–30, November/December 1999.

On Trust Establishment in Mobile *Ad-Hoc* Networks (Transcript of Discussion)

Laurent Eschenauer and Virgil D. Gligor

University of Maryland

Matt Blaze: This is a quick self-centered comment. The policies and credentials that you mention look like they would be quite readily implementable using KeyNote, particularly the different metrics of trust that you mention. Have you considered this?

Laurent Eschenauer: To use KeyNote, you need to collect the credentials, and that's really what we look at.

Virgil Gligor: We *are* considering implementing these things, or at least doing enough simulation to get a good grasp of what that implementation looks like. We believe that a fundamental area to work on would be this evaluation metric. At some point you have to commit to some metric. We have to look at what people have done in the past because we have to do on-line evaluation of evidence and the swarm-intelligence thing can actually help us to do this evaluation on line. We no longer need an infrastructure like the PKI, we build the infrastructure we need as we go along.

John Ioannidis: A good feature of KeyNote is that it's monotonic, so you cannot cause things not to happen by adding credentials.

Would there be any benefit in distributing all the information you receive, piggy-backed on the MANET AODV[1] routing algorithm? Are you doing that or have I missed the point?

Laurent Eschenauer: No. For the moment we are looking at maybe having an overlay network, so that the mobility does not impact me too much because the routing below me is going to take care of routing packets in the network. However, other people in the research team are looking at swarms (for example) for routing. It would be possible to combine those because both are an exploration of the network in the neighbourhood and it would diminish the overhead to do both at the same time. But for the moment we have been building the system as an overlay.

John Ioannidis: The routing also takes into consideration the appearance and disappearance of nodes.

Pekka Nikander: I was just wondering whether we have kind of chicken and egg problem here. If I understood correctly, you're trying to collect evidence for deciding whether you allow somebody to communicate with the rest of the network or not, and now you're using the same network for distributing that evidence as well.

[1] RFC 3561

Virgil Gligor: That was just one example, but it's a good example in the sense that we want it to do that. That's exactly the point.

Laurent Eschenauer: That is why, in the example we had, it is a server which is taking the load to continue requests with evidence to know if it's going to allow a request to pass in or out.

Pekka Nikander: Take a more complex situation like the US troops and the UK troops meeting at a battlefield and both of them have lost connections to the rest of the network. Both are suspicious because there should be the Rockies in between them, but they just happen to be together. How do you know that it's not suspicious, and so on?

Virgil Gligor: If you have an island and those guys were lost together and they had no access to the network and nobody sees them from the sky, then obviously they're isolated. In that case there is no sense in talking about access. I mean, that's certainly something that could happen, and life is tough. But what we are trying to do here is say that if someone somewhere generates evidence about them, then they establish this reachability.

To give you an example, if you find yourself in Singapore and you want to buy a coke with your PDA, and you are lost from the point of view of connectivity to your home domain. What do you do? When you entered the country, somebody stamped your passport - that's evidence. When you went to the hotel, they checked your passport and your credit card - that's evidence. So perhaps the hotel will authorise your PDA to buy a coke at the nearby store. That's the kind of thing.

Ross Anderson: I like the idea of using peer-to-peer networks as a basis, because the main problem with peer-to-peer networks is subversive insiders, and the main problem with battlefield equipment is battlefield capture; and these are in some sense the same thing. However, I'm very concerned at your implication that you can do quite fully distributed revocation. Suppose the rule is that someone may revoke, as suspected of capture, the equipment of someone of a lower rank. For example, a captain could revoke a lieutenant by saying "This lieutenant is now in the hands of the Iraqis, he is hereby revoked." The attack then goes as follows: a random captain is captured, the Iraqis can then send a message in respect of every lieutenant in the armed forces, saying "This lieutenant has now been captured." Thus the service denial attack is complete. The reason that you stop that normally is that there is a choke point through which the revocation has got to go such as the colonel in charge of the regiment that that particular lieutenant belongs to, and without that structure I don't see how you go forward.

Virgil Gligor: Here is the point I am trying to make. First of all, we did not address at all how revocations are authorised (the revocation policies). What we are concerned with here is simply the revocation mechanism which the kind of networks we are talking about, for example Freenet and swarm intelligence, use to direct revocation exactly where it is needed; there is no broadcasting, there are no certificate revocations that you push on anyone.

Ross Anderson: You're not answering my question, I'm saying, if random people are allowed to generate evidence that other random people are banned, there's no means of filtering that, and this becomes a mechanism for service denial attacks.

Virgil Gligor: Fake generation of evidence has to be taken into account by the evaluation metric, not by revocation. Revocation is a totally different topic. So what we address here is simply the mechanism, not whether or not a captain can revoke somebody's certificate, or anything like that. We have not addressed the policy aspect of revocation. So although your question is very relevant, it has not been addressed yet, we addressed only the mechanism for distributing revocation information fast. That is all we claim.

Laurent Eschenauer: I wanted to say something else. Your example will work if there is no single piece of other evidence about officer rank of the US, but there will be other pieces of evidence, by other people, saying that those guys are still alive, still good, still valid, and in service and doing very good work. So even if you capture one person and you use that person's credentials to generate fake evidence, it will only be evidence from one source. And that's why it's important in the metric to have independent paths of evidence and to build up the trust relationship using independent paths. That is why we need a system that distributes the most evidence to the most people so that even if you have one source of fake evidence, the rest of this good evidence is still available to everybody.

Ross Anderson: Yes, that's the problem. You are doing this in a distributed way if I, being a bad person, can call into question the honour of every soldier in the US army so that every other soldier in the US army has to start worrying about whether the guy on his left side is a traitor, and whether the guy on the right side is a spy. They would have to stop fighting for half an hour while they go and collect evidence to assure themselves that this guy's an American, and this guy's an American, and this guy's not an American, but he's a Brit so he's OK. I would have achieved a very remarkable result against your system. Now what I would suggest is that in a real life system you have to have some choke point for particular types of evidence. For example, in order to revoke a solider you must get a certificate signed by his colonel, and nobody is allowed to say bad things about the soldier except his colonel. That is how things tend to work in practice.

Laurent Eschenauer: I agree, that's in the metrics and the policy.

Virgil Gligor: There's one more thing that we didn't cover. We envisioned some of this network having some sort of a stable corner. We can harden these networks by providing a piece of the network that's almost always going to be alive and up and running. The area that we addressed with the example was the same area where we actually ran into these exceptional situations. So we haven't worked out the operational details. Clearly what you are saying is relevant, but we don't see that as being the Achilles' heel of this solution.

Bruce Christianson: What you're saying is, anybody can say anything about anybody else and you'll distribute that rapidly; whether you believe it or not is a policy decision.

Virgil Gligor: And whether we believe the evidence generated in this corner rather than evidence generated in that one.

Laurent Eschenauer: That was not the focus of our research. We really wanted to cover this evidence distribution issue.

Virgil Gligor: The trust metric is really the heart of this whole research.

Bruce Christianson: The trust metric isn't monotone. You might have said Ross is a bad guy, and I might have believed it. Then you might have said somebody else is a bad guy who I knew wasn't, so I stop believing you and I reinstate Ross.