Modeling Protocols for Secure Group Communication in Ad Hoc Networks (Extended Abstract)

Alec Yasinsac
Computer Science Department
Florida State University
yasinsac@cs.fsu.edu

Jim Davis
Department of Electrical and Computer Engineering
Iowa State University
davis@iastate.edu

Abstract

Since its introduction as a communications medium, wireless technology found broad application on the battlefield. Accompanying dramatic advances in wireless technology and the capabilities associated with small computing devices, the demand for advanced mechanisms to employ wireless technology in the battlefield continues to grow. We propose a model for deriving and reasoning about security protocols designed for battlefield use in ad hoc wireless networks. We contend that our model facilitates reasoning about protocols by integrating the communications and cryptographic aspects of battlefield group communication, and allows automated reasoning about resulting protocols. We illustrate our concept by introducing protocols to support special communication cases associated with the battlefield.

1. Introduction

The use of wireless networks is exploding as the limiting factors such as sufficient bandwidth, device size and weight, and power concerns are eliminated or mitigated. As a result, we are beginning to see the demand for small, highly mobile devices that utilize wireless communications to organize ad hoc networks that dynamically form, intercommunicate, and pass information to other wireless users and to wire-based networks, then dissolve. In this paper, we give a detailed description of one application of ad hoc wireless networks: battlefield communications. We provide a model that reflects the salient properties of the network and propose protocols that support this environment.

1.1. Communication on The Dynamic Battlefield [1]

The modern battlefield is highly dynamic. Units enter and leave and leave the battlefield continuously. The dynamic battlefield demands several characteristics of communications. Among them are that communication must be:

(1) **Fast**. While some limited setup may be tolerated before action starts, the ability to communicate during combat should be immediate in its accessibility to the transmitter and its delivery to the recipient.

(2) **Easy/transparent**. The transmitter must be able to communicate with minimal effort apart from their normal battlefield activities.

(3) **Available**. Parties must be able to communicate whenever they need to.

---

[1] We use the land battlefield example because it represents a rigorously dynamic and unpredictable environment for communications. Consequently, we utilize the battlefield terminology throughout this paper. Clearly, this topic parallels the field commonly referred to by terms such as "ad hoc wireless networks" or "dynamic mobile networks". We believe if our technology is effective in a battlefield environment, it will easily meet less rigorous environments such as combat ship-to-ship communications and a wide variety of industrial scenarios.

(4) **Authenticated**. The communication initiator must be able to absolutely identify all intended recipients.

(5) **Private**. The communications passed during combat should not be divulged to anyone not intended for receipt.

(6) **Integrity Protected**. Messages must be protected from modification during transmission.

(7) **Acknowledged**. All parties to the communication must know what the other parties did and did not receive.

We do not contend that this list is all-inclusive. Further, we recognize the impact of interactions of these requirements and posit that these interactions create the bulk of the complexity involved in secure battlefield communication.

1.2. Group Communication on the Battlefield.

Modern battlefield doctrine is based on mobility, flexibility, and rapid response to changing situations, yet also requires close coordination and mutually understood objectives among all members of the [command] group. This demands a group communication paradigm.

Group communication is sometimes thought of as broadcast technology. Broadcast and group communications are related, though not identical. Broadcast technology can provide efficient group communication, though group communication may or may not involve broadcasting messages.

The group communications paradigm is preferable over point-to-point connections in such an environment simply from the standpoint of reduced overhead. If the broadcast domains of the group members, the number of transmissions required to fully deliver a group message is minimized. If the broadcast domains are totally overlapping, group messages can be fully delivered with a single transmission.

While a broadcast medium enhances flexibility and efficiency, it also introduces security vulnerabilities. Since broadcast messages are available to anyone with a suitable receiver within the broadcast range, cryptography must be used to scramble messages for privacy and to authenticate intended recipients. For group communications, authentication and privacy are normally accomplished via security protocols and subsequent distribution of a group key.

1.3. Multicast Protocols to Support Group Membership

The reliance on group communication in tactical missions is critically important and a growing practice and research area. Multicasting is a popular mechanism for supporting group communication. In a multicast session, the sender transmits only one copy of each message that is replicated within the network and delivered to multiple recipients. The multicast group was developed from the concept of the *host group model* [CD85] in which "a host group is a set of network entities sharing a common identifying multicast address, all receiving any data packets addressed to this multicast address by senders (sources) that may or may not be members of the same group and have no knowledge of the groups' membership". Efficient bandwidth sharing is of paramount importance in low-bandwidth networks such as mobile ad-hoc networks. Although the host group definition allows non-members to send to the multicast address, secure multicasting designates that all senders must be authorized, whether or not they are also members of the group.

Multicasting in a wireless ad-hoc domain is significantly more complicated than multicasting in traditional wired networking. Multicasting in a wireless ad-hoc network behaves closer to multicasting across a LAN than multicasting across point-to-point links.

For ad-hoc networks, the scope of the multicasting can be divided into two major categories, namely intra-domain (within the ad-hoc network) and inter-domain (within the ad-hoc network and

beyond) multicasting. Essentially, the intra-domain case can be though of as a subset of the inter-domain multicasting case.

When multiple ad-hoc networks and/or wired networks are linked together in multicast groups, the problem adds several degrees of complexity. In such an environment, the majority of the complexity for mapping to the wired/existing infrastructure lies with the ad-hoc gateway. It is the responsibility of the gateway to map not only the multicasting connection itself between the ad-hoc network protocol and the wired protocol (PIM, DVMRP, etc.) but the gateway must also manage the security on the wired domain as well. Whereas the ad-hoc security mechanism would be known, the security mechanism for end-to-end service for the wired network may be dramatically heterogeneous in nature. Thus, the combination of these challenges (intra- and inter-domain) routing satisfying the security in multicast communication is a focus of this research.

1.4. Challenges to Managing Secure Groups

Not only are there numerous routing protocols, as seen in [Ra00, SGLA99], but there are many issues associated with managing a group. The group management issue becomes more complicated when the communication needs to be secure [Rei94]. Securing multicast communications involves distributing cryptographic keys to the members so that each can encrypt and decrypt messages as appropriate [HCD00, MRR99]. To maintain the security of encrypted packets, these keys must be recalculated and redistributed at designated times or upon certain events, such as when a member join or leave the group The group manager not only needs to be aware of membership changes but must also propagate the consequences of these membership changes to the rest of the group.

Numerous criteria are used to analyze secure multicast solutions [MRR99, CGIMNP99, SGLA99, WGL00]. These criteria are categorized into group membership management, network resource consumption, receiver resource requirements, sender resource requirements, and dependency upon particular standards. These categories are elucidated below.

- *Group membership management* criteria address the concerns of who is and is not part of the group, what the group looks like, and what happens if the group changes. Questions to consider when evaluating a solution's membership management capabilities are shown in Table 1.

- *Network resource consumption* criteria are concerned with the load on the network for various stages of the multicast communication process. When analyzing the bandwidth consumption of a solution, it is important to note how many messages must be transmitted each time a member joins or leaves and how large the control messages (those for managing the group) are in relation to the data messages. Also of importance is the volume of communication that can be effectively dealt with and whether the solution can handle bursty traffic.

- *Receiver and sender resource requirements* consider the following: How many keys must each member or sender store and how large are these keys? What is the processing time involved for the member or sender to, respectively, read or send messages? Does the solution allow non-members to send data? How many senders are allowed? Must these senders be known in advance of group creation?

- *Dependence upon standards* concern with whether the solution depends up a particular network protocol or network characteristics (such as stability, in order packet delivery, or reliable transmission)? Does the solution depend upon a particular application?

1.5. Key Distribution

Distributing keys in point-to-point environments is challenging. Keys must be generated and distributed in a way that guarantees that the security properties are upheld. One of the goals of this paper is describe a method and technology to generate and distribute group cryptography keys. A simple key management option for group keys is to distribute a single key to all members of the

group. One drawback of this method is that, since the key authenticates the group, a separate mechanism must be employed to prevent the possibility that the enemy may have acquired the group key. If a single group key is compromised, it is very difficult to detect.

1.6. Group Key Mechanisms

1.6.1 Diffie-Hellman key agreement

The foundation of many party-to-party key distribution schemes is based on the Diffie-Hellman public key scheme [DH76]. Several extensions to this technique have been proposed for group key distribution mechanisms [STW96, AST98]. The now well-known Diffie-Hellman computation relies on parties being able to raise large numbers to large powers under a large modulus. The idea is elegant in its conception and is widely considered the origin of public key cryptography.

Its elegance notwithstanding, the essence of the Diffie-Hellman computation has been somewhat more fleeting. Formal methods for protocol evaluation that easily model the effects of encryption and decryption, either mathematically [Mea91] or logically [BAN88], struggle to represent Diffie-Hellman in their formalism, though a recent breakthrough by Meadows [MN02] shows promise in this area.

1.6.2 Wright/Fischer [FW93]

Fischer and Wright developed a series of group key distribution protocols based on card games. Their protocols have the favorable property that they are effective even against adversaries with unlimited computing power. Unfortunately, their protocols are computationally intensive and not suitable for practical implementations.

1.6.3 Harn/Kiesler [HK89]

Harn and Kiesler offered a key distribution scheme using a Diffie-Hellman type computation among hierarchically organized groups with key distribution centers at the apexes. A favorable characteristic of their scheme is that each key distribution need maintain only a single secret key that is used to generate the shared key with each subordinate node. A detractor is that they require extensive setup and coordination so are not well suited to the dynamic group membership requirements of the modern battlefield.

2. A Model for Ad Hoc Battlefield Networks

We now present a model of communications components that facilitates discussions about battlefield group communication activities. We begin by defining the communicating nodes with four essential attributes: Identifier (ID), Most Recent Location (MRL), Transmission Range (XR), and Mobility Vector (MV). These attributes facilitate reasoning about immediate communication capabilities and allow nodes to predict connectivity in the dynamic environment. In our model, application of the attribute functions returns the attribute value of the entity (e.g. $ID(x)$ returns the identifier of the entity x).

We introduce the following functions to allow reasoning about the relationships of nodes, where x and y are communicating nodes.

Distance. $Dist(x,y)$ returns the distance between node locations

Broadcast Domain. $BD(x)$ returns the set of all nodes y where $dist(MRL(x),MRL(y)) < XR(x)$.

Communications Reach. $CR(x)$ returns the set of all nodes whose MRLs are within the XR of x and those that are within the communications reach of nodes in x's broadcast domain.

$$CR(x) = \{BD(x) \cup (CR(y) \text{ for all } y \in BD(x))\}$$

Full Partners. $FP(x)$ returns the set of all nodes y where $y \in CR(x) \wedge x \in CR(y)$

Transmitting Partners: $XP(x)$ returns the set of all nodes y where $x \in CR(y) \wedge y \notin CR(x)$.

Receiving Partners: RP(x) returns the set of all nodes y where $y \in CR(x) \wedge x \notin CR(y)$

2.1. Modeling Security Protocols for Ad Hoc Battlefield Networks

The essential operation performed in communications protocols is sending messages. Protocol outcome is determined by the data that the participants (and intruders) possess during and after these protocols. For our model, the functionality provided by the protocols requires a separate set of operators to those that model the physical communication medium.

We consider three set operations on messages: Computation (Mc), Possession (Mp), and Receive (Mr). Message Encryption (M_k) is reflected through subscript notation. Message possession is accomplished either by computing or receiving a message. Nodes possess messages that have been encrypted if they possess the ciphertext and the key.

$$x \in Mp(msg) <=> (x \in \{Mr(msg) \cup Mc(msg)\}) \vee (x \in \{Mp(msg_k) \cap Mp(k)\}$$

We modify the standard paradigm used to describe security protocols by allowing messages to be sent from one participant to a communications group. We can define a Communications Group as the set of all nodes that own the group key. Specifically,

CG(k) returns the set of all nodes that possess group key k.

Passing messages is the canonical activity in protocols. In our model, all messages are broadcast and may be connected to a message group. The message:

send(A,CG(k),msg)

depicts participant Alice sending a message to the communications group that shares key k. It is not clear from this specification what nodes receive the message. One may speculate that every node that holds key k receives the message, but this ignores communications limitations. We define message transmission by combining the network operators with the protocols operators and express the set of message recipients as:

for nodes x, y: $x \in Mr(msg)$ if $(sent(y,CG(k),msg) \wedge (x \in CG(k)) \wedge (x \in CR(y)))$

2.1.1 Modeling Group Diffie-Hellman

We now exercise the model to derive a protocol specific to the battlefield environment based on the variation of the Diffie-Hellman group key exchange protocol given in [AST98]. The setup requires each group member to establish a prior shared key with the group leader, C in this case.

$$Mp(kac) = \{C\}$$
$$Mp(kac^{-1}) = \{A\}$$
$$Mp(kbc) = \{C\}$$
$$Mp(kbc^{-1}) = \{B\}$$

The protocol messages occur sequentially, beginning with a non-group leader. All messages are transmitted in the clear, so no communications group is specified in the send operation.

$$send(A,,(g^{na}mod\ p)$$
$$send(B,,(g^{nb}mod\ p,g^{nanb}mod\ p)$$
$$send(C,,(g^{nbnckac}mod\ p,g^{nanckbc}mod\ p)$$

Once the messages are sent, it is possible to determine who has acquired the Diffie-Hellman key.

$$(B \in CR(A) \wedge C \in CR(B) \wedge \{A,B\} \subset CR(C)) => Mc(g^{nanbnc}mod\ p) = \{A,B,C\}$$

From this we can derive the result that:

$$CG(g^{nanbnc}mod\ p) = \{A,B,C\}$$

3. Modeling a Special Group Case with Silent Partners

On the modern battlefield, there may be groups that include passive members that receive from, but need not or cannot transmit to the group as a matter of noise or emissions discipline. In this case, there must be a protocol that allows selected members to participate in group key refreshment. These protocols must allow silent members to update their group key without participating and yet, remain in the group. We can represent this constraint in our model as follows:

$x \in SP(k)$ iff $(XR(x) = 0) \land x \in CG(k)$ for the group key k).

A simple solution for including a silent partner in key renewal is to employ a proxy that will act for them. Candidates for the task (potential proxies) are easily identified:

$x \in PP(y,CG(k))$ iff $(y \in SP(k)) \land y \in BD(x) \land \{x,y\} \subset CG(k)$).

The goal or our protocol is to generate a new key for the communication group of Alice, Bob, a Silent partner, and a proXy member:

Given: $CG(k) = \{A,B,X,S\}$

Derive: $CG(k') = \{A,B,X,S\}$

As with the earlier protocol, the setup requires each group member to establish a prior shared key with the group leader, this time B.

$Mp(kba) = \{B\}$

$Mp(kba^{-1}) = \{A\}$

$Mp(kbs) = \{B\}$

$Mp(kbs^{-1}) = \{S\}$

$Mp(kbx) = \{B\}$

$Mp(kbx^{-1}) = \{X\}$

In addition, the silent partner must establish a prior shared Diffie-Hellman value with the proxy .

$Mp(ns) = \{X,S\}$

The protocol proceeds sequentially, beginning with a non-group leader. The existing group key ensures authentication of the exchange.

$send(A,CG(k),(g^{na} \bmod p))$

$send(X,CG(k),(g^{nx} \bmod p, g^{ns} \bmod p, g^{nxna} \bmod p, g^{nsna} \bmod p, g^{nxns} \bmod p, g^{nsnxna} \bmod p)$

$send(B,CG(k),(g^{nbnsnxkba} \bmod p, g^{nbnanskbx} \bmod p, g^{nbnanxkbs} \bmod p)$

Once the messages are sent, it is possible to determine who has acquired the Diffie-Hellman key.

$(\{X,S,B\} \subset CR(A)) \land (\{A,S,B\} \subset CR(X)) \land (\{A,S,X\} \subset CR(B)) =>$

$Mc(g^{nanbnsnx} \bmod p) = \{A,B,S,X\}$

From this we can derive the result that:

$CG(g^{nanbnc} \bmod p) = \{A,B,S,X\}$

4. Reasoning About Subversion in Ad Hoc Networks

4.1. The Essence of Location

We use the term *subversion* to denote that an enemy has forcibly taken control of an asset, thereby diminishing the asset's capacity to carry out its tactical role. Within the context of communication equipment, there are several types of subversion, including: acquiring control from a distance (e.g., bogus control messages claiming to originate from a trusted source), inappropriate use of an asset due to subversion of the operator, and physical subversion (e.g., theft) that leads to an unexpected movement of the asset. For the purpose of this paper, we focus on the last problem of detecting unexpected movement of the asset.

Each asset in the battlefield has a role in supporting the tactical plan. Communication modalities, in so far as who an asset communicates with and for what reasons, are preplanned with rigorously defined procedures for deviating from the established plan. Movement of assets is also guided by the tactical plan. We desire to detect unexpected or unplanned movement of an asset, which may indicate subversion.

The model set forth in Section 2 allows us to reason about the movement of assets. The notion of proximity given there is of geographic location, such as a location defined by a measure such as latitude and longitude. We contend that the definition of location need not fit that traditional mold. Specifically, we may reason about the location of a node by considering the communication connections of each node and its neighbors. We refer to the former notion as geographic-based and the latter as path-based locality.

Our notion of path-based locality is founded the concept of broadcast domain, defined in Section 2. We consider nodes that are in the broadcast domain of another node to be that node's neighbors. Grouping neighbors of neighbors forms neighborhoods. For example,

Given: $BD(x) = \{q,r,s\}$ and $BD(y) = \{r,p\}$

Derive $NH(x,y) = BD(x) \cup BD(y) = \{q,r,s,p\}$

We may consider the Most Recent Location (MRL) of Section 2 to have a coarse granularity similar to that of a neighborhood, rather than exact spatial coordinates. Because of this, nodes need not transmit their coordinates via the ad hoc broadcast network. With regard to the subversion problem, an exact location may be less important as we desire to determine only if the node has left the neighborhood.

The difference in these two perspectives (geographic versus connection-oriented location) is seen in the rules used to model them. In the former perspective, the broadcast domain of a node x is defined to be the set of nodes whose location is within the transmission distance of x. In the latter, nodes are defined to be within the transmission distance of x if they are in x's broadcast domain. Fortunately, we easily model both of these locality notions and establish the strongest notion of location as:

$y \in BD(x)$ $(MRL(x),MRL(y)) < XR(x)$

When considering path-based locality, any node may infer its own location by knowing whom it can communicate with. During routine communications, a node will be able to passively determine which other nodes are in its proximity. A reasonable approximation to the present location can be constructed solely from a node's group peers rather than requiring an enumeration of all nodes within communication range.

Our strong definition of location of a node is illustrated by considering a designated authority to determine if the node is within the neighborhood. In a model-theoretic sense, the designated authority is defined to know the salient relationships between nodes. From a practical perspective, a monitor can detect node engagement in the secure group communications key agreement protocols, and can use that information to construct a set of nodes that that are able to communicate for each group formed.

Given a sufficiently large communication group, the self-location will be very close to the actual location. In our following discussion, we consider the global, model theoretic location to detect a subverted node.

4.2. Detecting Motion

The boundaries of a neighborhood are marked by the set of nodes in a group. There can be overlapping neighborhoods and nodes can simultaneously be members of many groups. The boundary for a neighborhood can be extended as a node moves in a particular direction, but at some point, it will move out of the communication range of the farthest node in the group. The group

becomes partitioned, and in the ensuing rekeying process, it is apparent that a node has moved. Over time, the mobile node will leave other groups as well and can be identified and located.

Consider the following example of a neighborhood {A,B,C,D,E,F,G}, where:

BD(A) = {A,B,C,D,F}
BD(B) = {A,B,C,D,F}
BD(C) = {A,B,C,D,F}
BD(D) = {A,B,C,D,E,F,G}
BD(E) = {D,E,F,G}
BD(F) = {A,B,C,D,E,F,G}
BD(G) = {D,E,F,G}

Following a motion, the set of nodes that X can communicate with changes. New nodes may come into communication range and distant nodes leave the set.

BD(A) = BD(A) - {B}
BD(B) = BD(B) - {C,D,F}
BD(C) = BD(C) - {B}
BD(D) = BD(D) - {B}
BD(E) = BD(E)
BD(F) = BD(F) - {B}
BD(G) = BD(G)

Because nodes move during the progression of the battle, this scenario may indicate a subversion only if the tactical battle plan does not call for the asset to relocate. Unexpected motion certainly raises concerns and may lower our confidence in the trustworthiness of the node.

Figure 1 shows our example of two neighborhoods (groups) of nodes. The boundaries of a neighborhood are marked by the set of nodes in a group. Here, nodes A, B, C, E, and F are in the same group and communicate through a common cryptographic key. Likewise, nodes E, D, F, and G are in a group and communicate by using a different key. The shaded circle center at node B represents the effective transmission range of B. In the right diagram, node B is moving away from its group. Note that B moves out of the range of group members C, D, and F thereby reducing its broadcast domain to node A.

Figure 1
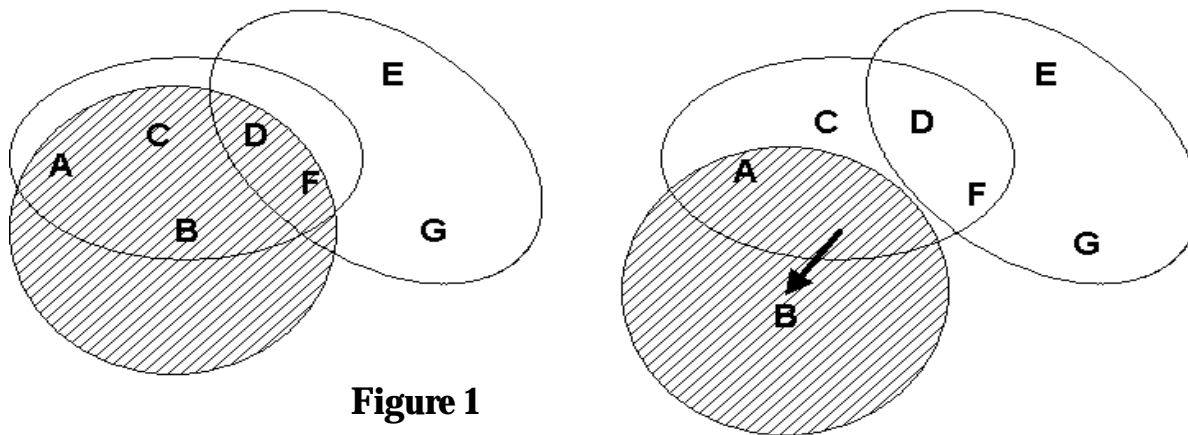
5. Conclusion

Ad hoc wireless networks are destined to be an essential element of the battlefield of the future, not to speak of the explosion of their use for personal and industrial applications. In this paper we present a model for reasoning about security characteristics of ad hoc wireless communication protocols. We exercised the model to illustrate a broadcast version of the well-known Diffie-Hellman group key distribution scheme and then developed a new protocol for a special case situation for battlefield networks employing Silent Partners. Finally, we showed how our model can be used to reason about the mobility of nodes and what this may say about the reliability or trust properties of those nodes.

6. Bibliography

[AST98] G. Ateniese, M. Steiner, and G. Tsudik, "Authenticated Group Key Agreement and Friends", In ACM CCS ACM, November 1998

[BAN88] Burrows, M., Abadi, M., and Needham, R. M. "A Practical Study in Belief and Action", In Proc of the 2nd Conf on Theoretical Aspects of Reasoning about Knowledge (Asilomar, Ca., Feb. 1988) M. Vardi, Ed. Morgan Kaufmann, Los Altos, Calif., 1988, pp. 325-342

[CD85] D.R. Cheriton and S.E. Deering. "Host Groups: A Multicast Extension for Datagram Internetworks". In 9th Data Communication Symposium, September 1985

[CGIMNP99] R. Canetti, J. Garay, G. Itkis, D. Micciancio, M. Naor, and B. Pinkas. "Multicast security: A taxonomy and some efficient constructions", In *Proc. IEEE Infocom*, March 1999

[DH76] W. Diffie and M. Hellman, "New Directions In Cryptography", *IEEE Transactions on Information Theory*, IT-22(6):644-654, November 1976

[DOW92] W. Diffie, P. C. van Oorshot, and M. J. Wiener, "Authentication and Authenticated Key Exchanges", *Designs, Codes and Cryptography*, 2(2):107-125, June 1992

[FW93] Michael Fischer and Rebecca N. Wright, "An Efficient Protocol for Unconditionally Secure Secret Key Exchange", Proc of the 4th Symp on Discrete Algorithms (1993), pp. 475-83

[HCD00] T. Hardjono, B. Cain, and N. Dorawswamy, "A framework for group key management for multicast security", IETF Internet draft, August 2000. draft-ietf-ipsec-gkmframework-03.txt

[HK89] Lein Harn & Thomas Kiesler, "Authenticated Group Key Distribution Scheme For a Large Distributed Network", IEEE Symposium on Security and Privacy, 1989, pp 300-309

[Mea91] Meadows, C., 'A System for the Specification and Analysis of Key Management Protocols'. From 1991 IEEE Computer Society Symp on Research in Security and Privacy, pp. 182-195.

[MN02] C. Meadows and P. Narendran, "A Unification Algorithm for the Group Diffie-Hellman Protocol", WITS (in conjunction with POPL'02), Portland, Oregon, USA, Jan 14-15, 2002

[MRR99] M. J. Moyer, J. R. Rao, and P. Rohatgi, "A survey of security issues in multicast communications". *IEEE Network*, pp. 12-23, Nov/Dec 1999

[NS78] Roger M. Needham, Michael D. Schroeder, "Using Encryption for Authentication in Large Networks of Computers", CACM, December 1978 vol. 21 #12, pp. 993-999

[Ra00] M. Ramalho, "Intra and inter-domain multicast routing protocols: A survey and a taxonomy", *IEEE Communications Surveys & Tutorials*, vol. 3, no. 1, pp. 2-25, First Quarter 2000

[Rei94] M. Reiter, "A secure group membership protocol", Proceedings of the IEEE Symposium on Research in Security and Privacy, Oakland, CA, May 1994

[SGLA99] C. Shields and J.J. Garcia-Luna-Aceves, "KHIP -- A scalable protocol for secure multicast routing", *ACM SIGCOMM*, 1999

[STW96] M. Steiner, G. Tsudik, M. Waidner, "Diffie-Hellman key distribution extended to group communication", In Proc. 3rd ACM CCS, New Dehli, India, 14-16 May, 1996, pp. 31-7

[WGL00] C. K. Wong, M. Gouda, and S. S. Lam, "Secure group communications using key graphs", *IEEE/ACM Transactions on Networking* vol. 8, no. 1, pp. 16-30, February 2000