

Authentication for Mobile IPv6

Anupam Datta John C. Mitchell Frederic Muller
Dept. Computer Science, Stanford University
<http://www.stanford.edu/~danupam,~jcm/>

Dusko Pavlovic
Kestrel Institute, Pale Alto, CA
<http://www.kestrel.edu/>

Abstract

We present a protocol for authenticating Mobile IPv6 connections, addressing requirements established in the relevant Internet Draft [2]. The protocol imposes minimal computational requirements on mobile nodes, uses as few messages as possible, and may be adapted to resist denial of service attacks. Our protocol has two parts, an initialization phase and an update phase. The initialization phase can take advantage of local public-key infrastructure, a bidirectional chain of trust, or operate without prior authenticating information. Each execution of the update phase uses the shared secret established in the previous phase to maintain security of the mobile connection. We have formally verified the correctness of the protocol using the finite-state analysis tool Mur ϕ , which has been used previously to analyze hardware designs and security properties of several protocols.

1. Introduction

In Mobile IPv6 [1], a mobile node has two associated IP addresses. A mobile node is always identified by its *home address*. While operating away from its home, a mobile node is also associated with a *care-of address*, which provides information about its current location on the internet. The care-of address is registered with the home agent, which transparently routes IPv6 packets sent to the home address to the care-of address. To reduce routing distance and relieve the load on home agents, a mobile node may also inform other IPv6 nodes about its current care-of address by sending a *binding update*. The need for authenticating binding updates has been recognized [1, 2]. However, previous authentication proposals either fall short of meeting the authentication requirements [3, 4] or require the mobile nodes to perform expensive public key operations.

In this paper, we present a protocol for authenticating binding updates. The protocol has two phases: (a) a once-per-connection *initialization phase* in which two mobile nodes use any available chain of trust through their home

agents to confirm a shared secret; (b) an *update phase* that is executed every time an authenticated binding update needs to be sent. The protocol provides sender authentication, data integrity, and replay protection, without requiring mobile nodes to perform public key operations. Using a forward message from the mobile node to the corresponding node, and a reply through the home agents only once in the initialization phase, the protocol minimizes the number of messages exchanged between participating entities. While we assume that each mobile node shares a secret key with its home agent, the protocol does not require home agents to maintain state during the execution of the protocol, avoiding memory denial-of-service attacks and other resource consumption problems.

The most difficult adoption issue for any form of authenticated Mobile IP is reliance on authentication infrastructure. Our goal is to make the best use of whatever infrastructure is available, in no case offering less assurance than the link between the two home agents, which in some cases may be wired, unauthenticated IP. As an expository convenience, we present our protocol under the assumption that each home agent has a public-private key pair, with each public key known to the other. We then consider three other scenarios and present variations of the initialization phase that are appropriate to each scenario. The update phase of our protocol remains the same in each case, since the basis for authenticated update is the shared private information established by the initialization phase. While there have been concrete proposals for maintaining a global trust infrastructure, e.g., by extending the Domain Name System (see [5, 6]), we believe our protocol offers significant advantages over previous proposals in the absence of public-key infrastructure. Perhaps the most likely configuration in the near term is an SSH-like scheme, using unauthenticated key exchange between home agents the first time these home agents participate in Mobile IP, and caching the shared keys for future use by any mobile nodes using the same pair of home agents. While this configuration is susceptible to a person-in-the-middle attack, an attack is only possible once for each pair of home agents (up to the capacity of the key cache).

We have formally verified the correctness of our protocol using the finite state analysis tool $\text{Mur}\varphi$ [7]. While formal methods have been successfully used to analyze key exchange and authentication protocols [8, 9, 10, 11, 13, 12], this is the first Diffie-Hellman based key exchange protocol that has been analyzed with $\text{Mur}\varphi$. In addition, previously analyzed protocols such as SSL [14] and contract signing [15, 16] involve only two parties while Mobile IP has four (two mobile nodes and their associated home agents). Therefore, the analysis of Mobile IP requires several new concepts and modelling techniques.

The remainder of this paper is structured as follows. Section 2 describes the requirements for security in Mobile IPv6. Section 3 briefly discusses the previous proposals for authenticating binding updates. In Section 4, we present our basic protocol. Section 5 presents our modelling assumptions and analysis results. In Section 6, we discuss extensions to prevent denial of service attacks. Section 7 describes how our base protocol can be adapted to work in a AAA infrastructure [17, 18] and weaker trust infrastructures. Concluding remarks appear in Section 8.

2. Security Requirements for Binding Updates

While the binding update feature of Mobile IPv6 makes end-to-end routes shorter, the ability to change routes on the fly introduces a number of security concerns. A Mobile IPv6 binding update specifies an association between the home address of a mobile node and its a care-of address, along with the remaining lifetime of that association. Upon receiving a binding update from a mobile node (MN), a correspondent node (CN) creates a binding cache entry and stores this association. Subsequently, the CN will send all packets destined for the MN to its care-of address instead of its home address. Threats and security requirements for binding update are described at length in [2]. Here, we discuss the most pressing security issues.

In the absence of a pre-established security association between MN-CN pairs, two major security threats arise:

- An attacker may tamper with the binding cache entries since binding updates will not be authenticated. In fact, an active attacker can launch a person-in-the-middle attack, becoming the default router to the MN and from the MN to the CN.
- An attacker can launch denial-of-service attacks on MN's, CN's and Home Agents (HA's). This could be done, for example, by flooding IPv6 nodes with fake binding updates.

Our main goal is to institute a mechanism for setting up security associations between MN-CN pairs that will allow binding updates to be authenticated. Additional requirements, as specified in [2], are:

- Identity verification should not rely on the existence of a global PKI.
- Minimize the number of messages and bytes sent between the participating entities.
- Consider the computational capabilities of the MN's and CN's.
- Resist denial-of-service attacks.
- In any event, provide no weaker guarantees than IPv4.

The starting point for our investigation of authenticated binding updates is the realization that every authentication protocol relies on some form of previously established shared information. In general, an authentication protocol involves a *claimant* A and a *verifier* B . Presented with a purported identity of the claimant, the goal of the verifier is to verify the claimed identity. In the setting of Mobile IP, the IP address of an agent can be taken as its identity.

In order for claimant A to directly authenticate herself to verifier B , agent A must be able to perform some action, either individually or in collaboration with services available on the network, that can only be performed by A and that B is able to recognize, either individually or in collaboration with services available on the network. If A has a public-private key pair and the public verification key is known by B to be associated with A , then digital signatures allow B to verify a claim from A . Shared symmetric keys, shared secret hash keys, and other shared secret or unforgeable information can also be used for authentication. However, unless there is some kind of secret information, known to A and not known to a potential impersonator, and verifier B has prior knowledge that allows B to recognize some kind of operation involving the secret information held by A , there is no known way for B to reliably authenticate A .

Since any protocol that truly authenticates binding updates will depend on some form of chain of trust, our goal is to present the most reliable form of authentication possible for each of the most realistic or potentially achievable chains of trust. Further discussion on how our authentication protocol leverages various trust infrastructures appears in Section 7.

3. Previous Proposals

Two previous protocols for authenticating binding updates [3, 4] are susceptible to person-in-the-middle attacks. We review these protocols to show some of the forms of cryptography that have been considered and to illustrate their shortcomings.

Bradner, Mankin and Schiller [3] proposed a framework called *Purpose Built Keys*. Before initiating a connection

with a correspondent node (CN), a mobile node (MN) generates a public-private key pair (called a purpose-built key pair) for use during the connection. The MN then sends a hash of the public key to CN. Subsequently, when MN sends a binding update, MN signs it with its private key and also sends the public key to CN. CN verifies that the public key hashes to the same value it had received before and, if so, uses the public key to verify the digital signature. An advantage of this framework is that it does not require any security infrastructure. However, purpose-built keys provide authentication if and only if the initial hash of the public key is received correctly by CN. This might not be the case. An attacker could intercept the hash and send the hash of a different key (which it owns) to CN. Subsequently, it can pretend to be MN without CN being any wiser. The authors of the draft were, of course, aware of this weakness.

Le and Faccin [4] propose two protocols for authenticating binding updates. The first assumes that both the MN and the CN share security associations with two AAA servers (e.g., RADIUS [17], DIAMETER [18]) and these two servers in turn share a security association. The protocol then uses this chain of trust to achieve authentication. Our protocol uses a similar architecture and can be easily adapted to work within a AAA infrastructure. Furthermore, it employs fewer messages than Le and Faccin’s protocol (5 as opposed to 7). This improvement is obtained by combining encryption-based and signature-based authentication techniques and will be elaborated further in the next section. The second protocol proposed in [4] involves an unauthenticated Diffie-Hellman key exchange between MN and CN. The resulting key is subsequently used to authenticate binding updates. The authors recognize that this protocol is vulnerable to a man-in-the-middle (MITM) attack but state that “due to the properties of IP” such an attack will always be detected. They argue that since an attacker cannot remove any packets from the network, if a MITM attack is launched, both the MN and the CN will receive two Diffie-Hellman exponentials and will therefore be able to detect the attack. There are a couple of objections to this argument. Firstly, an attacker could remove packets from the network, e.g., if it gained control of the router being used by MN. Secondly, even if it were not able to remove packets from the network, it could delay them, e.g., by flooding MN and CN’s network buffers with junk packets. During that interval, it could potentially do a lot of damage. Also, even if MN and CN detect the attack, the only thing they can do is drop the session. So, the attack still succeeds as a denial of service attack.

4. The Protocol

In this section, we present the authentication protocol based on our strongest assumptions. We assume that each

mobile node (MN) has a pre-established bidirectional security association with its home agent (HA) using which they can authenticate each other. This is a reasonable assumption (and has been recognized as such in [2]) since typically the MN and its HA will belong to the same administrative domain. We also assume that HA’s are capable of authenticating each other using public key cryptography.

4.1 Notation

The following notation is used in describing the protocol.

MN	Mobile node
CN	Correspondent node
HA_{MN}	Home agent of MN
HA_{CN}	Home agent of CN
BU	Binding update
K_M	Shared secret between MN and HA_{MN}
K_C	Shared secret between CN and HA_{CN}
K_M^+	Public encryption key of HA_{MN}
$Cert_i$	Certificate of entity i
$Sign_i\{\dots\}$	Signed by entity i
$\{\dots\}_K$	Encrypted with key K
$[x]$	x is an optional parameter
IP_i	IP (home) address of i
K_D	Shared Diffie-Hellman secret key
$h(k, M)$	Keyed hash of message M with key k
s_{MN}	$\{IP_{MN}, IP_{CN}, g^x, g^y\}_{K_D}$
s_{CN}	$\{IP_{CN}, IP_{MN}, g^y, g^x\}_{K_D}$

4.2 Protocol Description

Our protocol consists of two phases: (a) an *initialization phase* in which MN and CN set up an authentication key; (b) an *update phase* in which MN sends an authenticated binding update to CN using the key obtained from phase (a).

4.2.1 Initialization Phase

The initialization phase is an authenticated Diffie-Hellman key exchange protocol [19] resulting in a shared secret between an MN and a CN . The protocol is shown in Figure 1. We first take a closer look at the protocol and discuss the purpose served by the different message components. Then, we describe some of the features of this protocol.

Message 1 is straightforward. Note that it assumes that MN already knows a group and generator that is acceptable to CN .

Upon receiving message 1, CN generates g^y , computes $K_D = g^{xy}$ and s_{CN} and sends a message to HA_{CN} encrypted with their shared key, K_C . Note that the use of encryption-based authentication ensures that no attacker

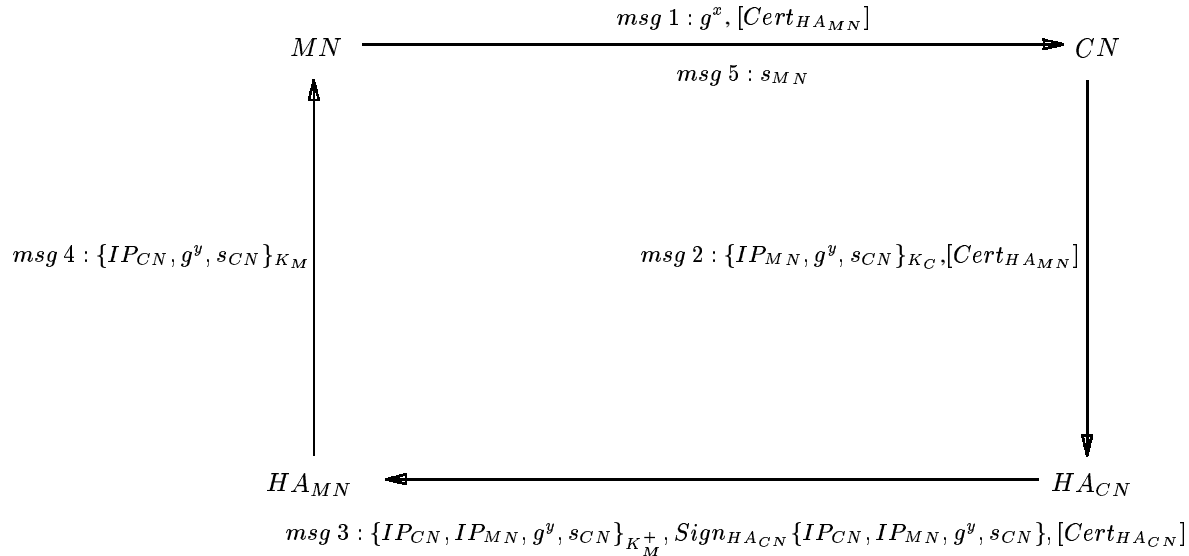


Figure 1. Initialization Phase

can observe g^y . s_{CN} consists of the session parameters encrypted with the session key, K_D . It is forwarded by the two home agents and proves to MN that CN has computed the same secret key. Thus, the protocol provides *direct authentication*, a desirable property as described in [21]. If direct authentication is not required, then s_{CN} can be replaced by g^x in messages 2, 3 and 4 resulting in a lighter-weight protocol.

In Message 3, H_{ACN} forwards to H_{AMN} the session parameters that it recovers by decrypting message 2. It is necessary to include CN 's IP address since the same MN might be executing this protocol parallelly with multiple correspondent nodes. Encrypting the session parameters with public key of H_{AMN} ensures that no other entity can recover g^y . The signature proves to H_{AMN} that the message originated from H_{ACN} . Note that the signature scheme should be such that no information regarding plaintext data can be deduced from the signature itself on that data (e.g., when the signature operation involves preliminary one-way hashing). This is critical because, in general, data may be recovered from a signature on it (e.g., RSA without hashing). An alternative approach would be to include the signed block inside the public key encryption. A disadvantage of this method is that here the data to be public-key encrypted is larger. This might require adjustment of the block-size of the public encryption scheme or the use of techniques like cipher-block-chaining (see Section 12.5.2 of [20] for further discussion on the relative advantages of the two methods).

H_{AMN} forwards the session parameters to MN in message 4, encrypted with K_M , their shared key. Encryption serves the same purpose as before: it hides g^y from all other entities. We note here that the structure of the plaintext encrypted with K_M in this message is quite similar to that encrypted with K_C in message 2. While this does not lead to any attack under the assumption of black box cryptography, it might be useful to break the symmetry by reordering two message components, e.g., g^y and s_{CN} in message 4. It has been recognized that symmetries in a protocol should be used with caution, due to both the possibility of reflection attacks, and attacks in which responses from one party can be reused within a protocol (see [21]).

In message 5, MN encrypts the session parameters with the newly computed Diffie-Hellman secret. This serves two purposes. Firstly, it proves to CN that the message actually came from MN (since the only entities which know g^y are CN , H_{ACN} , H_{AMN} , and MN and we assume that these entities are honest). Secondly, it provides proof that both MN and CN have computed the same secret, i.e., it provides direct authentication.

In this protocol, CN authenticates herself to MN through the chain of trust $CN \rightarrow H_{ACN} \rightarrow H_{AMN} \rightarrow MN$. MN could also use the same trust chain in the other direction to authenticate herself to CN . This would lead to a 7-message protocol similar to the one presented in [4]. In order to reduce the number of messages to 5, we use a different technique. CN 's Diffie-Hellman exponential is never

sent in the clear; it is sent encrypted along the trust chain and the very fact that MN is able to recover it and compute the same Diffie-Hellman secret as CN serves to verify her identity. We thus combine the two standard techniques of signature-based and encryption-based authentication to realise a protocol with fewer messages. The difference in the number of messages exchanged in the two protocols becomes more appreciable as the trust chain becomes longer. If there are n participants, our protocol requires $n + 1$ messages while the previous approach requires $2n - 1$. It also appears that we can do no better. Specifically, any Diffie-Hellman based authenticated key exchange protocol P between entities A and B that uses a chain of trust of length n and provides direct authentication requires at least $n + 1$ messages.

The use of Diffie-Hellman key exchange ensures that the protocol provides *perfect forward secrecy* (PFS), i.e., the disclosure of long-term secrets like private signing/decryption keys does not compromise the secrecy of exchanged keys from earlier runs. However, if perfect forward secrecy is not a requirement, an alternative would be to replace the exchange of DH exponentials by an exchange of nonces (n_{MN} and n_{CN}). n_{MN} could be sent in the clear in message 1, while n_{CN} is sent encrypted in message 2, 3, 4. The shared secret could be derived from a hash of these nonces, e.g., $\text{HMAC}(n_{MN}, n_{CN})$ (see [22]). Note that this variation of the base protocol does not have PFS since the compromise of long term secrets like the shared key between CN and HA_{CN} exposes all past n_{CN} values and hence all past session keys. The advantage is that it is computationally very lightweight: mobile nodes have to perform relatively inexpensive operations - generating a random nonce and computing a hash instead of exponentiation.

Typically, in order to prevent replay attacks, key exchange protocols require each participant to use some fresh information in every run of the protocol. The Diffie-Hellman exponentials serve this purpose in our protocol. We would also like to note here that an implicit assumption in our protocol is that it is possible to verify whether a home agent actually owns a mobile node which it claims as its own (e.g., by matching the network prefix of the home address of the mobile node with that from the home agent's certificate). Otherwise, the protocol is open to attack. Any adversary which possesses a certificate could intercept message 1 and then send message 3 without HA_{MN} being any wiser.

4.2.2 Update Phase

Once MN and CN have set up a shared secret, K_D , MN can easily send an authenticated binding update (BU) by executing the following 1-message protocol.

$$BU, h(K_D, BU)$$

Here, $h(\dots)$ is a keyed cryptographic hash function (e.g., HMAC [22]). This protocol provides sender authentication since MN is the only entity other than CN which possesses the key K_D . Data integrity is also provided since the hash is also a message authentication code (MAC). Replay attacks can be prevented by using the *sequence number* field in the binding update option (see Section 5.1 of [1]). The sequence number field holds an 8-bit number. Each binding update sent by a mobile node must use a sequence number which is greater than the sequence number of the previous binding update sent to the same destination address. Every time the sequence number space is exhausted, the shared key should be refreshed. Key refresh could, of course, be done by re-executing the initialization phase protocol. Alternatively, MN and CN could set up a new key by executing an authenticated Diffie-Hellman key exchange protocol in which they can use the old key to authenticate messages. This approach would require the home agents to be involved in only the initialization phase (once per (MN, CN) pair) and the system would scale up better.

5. Analysis of the Protocol

We used $\text{Mur}\varphi$, a finite-state analysis tool, to carry out a formal analysis of the initialization phase of our protocol. In this section, we briefly outline the general methodology and describe some of the challenges we faced in applying it to this protocol. A detailed description of the methodology can be found in [13].

5.1 The $\text{Mur}\varphi$ Verification System

$\text{Mur}\varphi$ [7] is a finite-state verification tool that has been successfully applied to multiprocessor cache coherence protocols and multiprocessor memory models [23, 24]. The purpose of finite-state analysis (also called *model checking*) is to exhaustively search all possible execution sequences. While this process often reveals errors, failure to find errors does not imply that the protocol is completely correct, because the $\text{Mur}\varphi$ model may simplify certain details and is inherently limited to configurations involving a small number of protocol participants.

To use $\text{Mur}\varphi$ for verification, one has to model the protocol in the $\text{Mur}\varphi$ language and augment this model with a specification of the desired properties. The $\text{Mur}\varphi$ system automatically checks, by explicit state enumeration, if all reachable states of the model satisfy the given specification. For the state enumeration, either breadth-first or depth-first search can be selected. Reached states are stored in a hash

table to avoid redundant work when a state is revisited. The memory available for this hash table typically determines the largest tractable problem.

The Mur φ language is a simple high-level language for describing non-deterministic finite-state machines. Many features of the language are familiar from conventional programming languages. The main features not found in typical high-level programming languages are described in the following paragraphs.

The *state* of the model consists of the values of all global variables. In the *startstate* statement, initial values are assigned to global variables. The transition from one state to another is performed by *rules*. Each rule has a Boolean condition and an action, which is a program segment that is executed atomically. The action may be executed if the condition is true (i.e., the rule is enabled) and typically changes global variables, yielding a new state. Most Mur φ models are nondeterministic since states typically allow execution of more than one rule. For example, in the model of our authentication protocol, the intruder (which is part of the model) usually has the choice of several messages to replay.

Mur φ has no explicit notion of *processes*. Nevertheless a process can be implicitly modeled by a set of related rules. The *parallel composition* of two processes is simply done by taking the union of the rules of the two processes. Each process can take any number of steps (actions) between the steps of the other. The resulting model is that of *asynchronous, interleaving concurrency*.

The Mur φ language supports *scalable* models. In a scalable model, one is able to change the size of the model by simply changing constant declarations. When developing protocols, one typically starts with a small protocol configuration. Once this configuration is correct, one gradually increases the protocol size to the largest value that still allows verification to complete. In many cases, an error in the general (possibly infinite state) protocol will also show up in a down-scaled (finite state) version of the protocol. Mur φ can only guarantee correctness of the down-scaled version of the protocol, but not that of the general protocol. For example, in modelling our authentication protocol, the numbers of mobile nodes and home agents were scalable and defined by constants.

The desired properties of a protocol can be specified in Mur φ by *invariants*, which are boolean conditions that have to be true in every reachable state. If a state is reached in which some invariant is violated, Mur φ prints an error trace - a sequence of states from the start state to the state exhibiting the problem.

5.2 The Methodology

We analyzed our protocol using the following sequence of steps:

1. *Formulate the protocol.* This generally involves simplifying the protocol by identifying the key steps and primitives. The Mur φ formulation of a protocol, however, is more detailed than the high-level descriptions often seen in the literature, since one has to decide exactly which messages will be accepted by each participant in the protocol. Since Mur φ communication is based on shared variables, it is also necessary to define an explicit message format, as a Mur φ type.
 2. *Add an adversary to the system.* We assume that the adversary (or intruder) can masquerade as an honest participant in the system, capable of initiating communication with a truly honest participant, for example. We also assume that the network is under the control of the adversary and allow the adversary the following actions:
 - overhear every message, remember all parts of each message, and decrypt ciphertext when it has the key;
 - intercept (delete) messages;
 - generate messages using any combination of initial knowledge about the system and parts of overheard messages.
- Although it is simplest to formulate an adversary that nondeterministically chooses between all possible actions at every step of the protocol, it is more efficient to reduce the choices to those that actually have a chance of affecting other participants.
3. *State the desired correctness conditions.* A typical correctness condition would be that the intruder does not learn any secret information. More details about the correctness conditions for our protocol are given in Section 5.4.
 4. *Run the protocol* for some specific choice of system size parameters. We have been able to run our protocol with upto 3 mobile nodes and 3 home agents, where each mobile node can execute 2 sessions in parallel. Details of execution time appear in Section 5.4.

5.3 The Intruder Model

The Mur φ intruder model is limited in its capabilities and does not have all the power that a real-life intruder may have. In particular:

- *No cryptanalysis.* Our intruder ignores both computational and number-theoretic properties of cryptographic functions. As a result, it cannot perform any cryptanalysis whatsoever. If it has the proper key, it can read an encrypted message or (forge a signature).

Otherwise, the only action it can perform is to store the message for a later replay.

- *No probabilities.* Mur φ has no notion of probability. Therefore, we do not model “propagation” of attack probabilities through our finite-state system (e.g., how the probabilities of breaking the encryption, forging the signature, etc. accumulate as the protocol progresses). We also ignore, e.g., that the intruder may learn some probabilistic information about the participants’ keys by observing multiple runs of the protocol.
- *No partial information.* Keys, signatures, etc. are treated as atomic entities in our model. Our intruder cannot break such data into separate bits. It also cannot perform an attack that results in the partial recovery of a secret (e.g., half of the secret bits).

In spite of the above limitations, previous studies have shown that Mur φ is a useful tool for analyzing security protocols. It considers the protocol at a high level and helps discover a certain class of bugs that do not involve attacks on cryptographic functions employed in the protocol. Mur φ is quite useful in discovering authentication bugs since properties like key ownership, source of messages, etc. are easily captured in logical statements. The fact that Mur φ did not uncover any bugs in our protocol therefore gives us a fair degree of confidence in its correctness.

5.4 Modelling the Initialization Phase

The Mur φ model of the protocol consists of three types of finite state machines corresponding to mobile nodes, home agents and intruders. The number of mobile nodes, home agents and intruders are scalable and defined by constants. Each mobile node can participate in a number of parallel sessions. This number can also be configured by changing the value of a constant.

The state of a mobile node consists of her IP address, the IP address of her home agent, the secret she shares with her home agent and the individual states of all the sessions that she is involved in. We associate the *state-id* of a session with the next message that the node is expecting in that session. We denote by S_i the state in which a node is expecting the i^{th} message of the protocol. For example, after initiating a session (sending message 1), a node sets the state-id of that session to S_4 since the next message that it expects is the 4th message of the protocol. Initially, the state-ids of all sessions are set to S_1 . In this state, a mobile node can spontaneously initiate a session. Other than the state-id, the state of a session also includes the values of all the session parameters that the mobile node has seen up to that point. For example, if the state-id of a session that a mobile node is taking part in is S_4 , then the state would also contain the

node’s Diffie-Hellman private key and the address of the peer node with whom she is executing the session. Upon completing a session, the state-id for that session is set to S_{DONE} . The transition rules for a mobile node capture the exact sequence of actions that she would carry out in an actual run of the protocol. For example, when a mobile node receives the 4th message of the protocol, she processes it iff the corresponding state-id is S_4 . She then verifies that the encryption key specified in the message is the same as the key she shares with her home agent. This corresponds to decrypting the message. Then she computes the Diffie-Hellman secret using the Diffie-Hellman exponential in the message and her own previously recorded Diffie-Hellman private key. She finally verifies that the computed secret matches the one in the message before changing the state-id to S_{DONE} .

The finite state machine of a home agent is much simpler. Since a home agent only forwards messages, his state does not change during the execution of the protocol. The state of a home agent consists of his certificate, the verification key of the trusted third party (Certification Authority of the PKI) and the shared keys with the mobile nodes in his domain. The transition rules define the sequence of actions that a home agent executes upon receiving the 2nd or 3rd message of the protocol.

As mentioned before, the intruder’s transition rules enable it to intercept messages, overhear all messages and remember parts of all overheard messages and generate new messages using any combination of initial knowledge and parts of overheard messages. Thus, at any given point in the protocol, there are a large number of possible transitions for the intruder. This results in a very large number of reachable states for the protocol. The following techniques proved useful in reducing the number of states to be explored:

- The intruder always intercepts all messages sent by the honest participants.
- The intruder does not send messages to honest participants in states where at least one of the honest participants is able to send a message.
- The intruder only generates messages that are expected by the legitimate parties and that can be meaningfully interpreted by them in their current state, e.g., the intruder sends the 4th message to a mobile node only if the mobile node is in state S_4 .

The first two techniques have been proved to be sound (see [25]), i.e., each protocol error that would have been discovered in the original state graph will still be discovered in the reduced state graph. The soundness of the third technique is quite obvious.

We modelled the following correctness conditions in our Mur ϕ code:

- If two mobile nodes have completed a session with each other, then the shared Diffie-Hellman secret computed by both must be identical.
- The secret shared between two mobile nodes is not in the intruder’s database of known message components.
- The mobile nodes agree on each other’s identity and protocol completion status, i.e., if A has completed a session with B , then B should also have completed the same session with A or B should be waiting for the 5th message of the protocol.

Mur ϕ did not discover any violations of the above mentioned correctness conditions for the configurations on which we ran the verifier. Running under Linux on a 300 MHz dual-processor Pentium II with 512MB of RAM, the verifier required approximately 30 seconds to check for the case with 2 mobile nodes, 2 home agents and no more than 2 simultaneous sessions per mobile node. About 4200 states were explored. The largest instance of our model that we verified included 3 mobile nodes, 3 home agents and no more than 2 simultaneous sessions per mobile node. Checking took about 20 minutes, with 125,941 states explored.

We note here that this is the first Diffie-Hellman key exchange based protocol that has been modelled using Mur ϕ . Modelling the DH-protocol within the finite state analysis framework required some thought. The “obvious” approach would be to do what is actually done in practice: use integers to explicitly model the various parameters (g , p , x , y) and then derive the secret by exponentiating modulo p . However, we observe that explicit exponentiation is unnecessary. The two main properties that the model should capture are: (a) the computational hardness of the Diffie-Hellman problem, i.e., given $g^x \bmod p$ and $g^y \bmod p$, it should not be possible to compute $g^{xy} \bmod p$; (b) the commutativity of exponentiation so that $(g^x \bmod p)^y \bmod p = (g^y \bmod p)^x$. (From Fermat’s Little Theorem, we also know that both these values are equal to $g^{xy \bmod (p-1)} \bmod p$. In order to capture these two properties we used a different technique. In the Mur ϕ code, a *key* is modelled as a record with two fields - a *type* and a *random value*. We defined three additional types of keys: *dh-private*, *dh-public*, and *dh-secret*. x is of type *dh-private* and g^x is of type *dh-public*. For both keys, the value is x . The DH secret is computed using a function which takes in a key of type *dh-private* (say with value x) and another of type *dh-public* (say with value y) and returns a key with type *dh-secret* and value $xy \bmod (p-1)$. Note that interchanging the values of x and y yields the same secret because of the commutativity of integer multiplication. This ensures that both participants compute the same Diffie-Hellman secret. Also, since this

function is the only way of computing a DH secret, property (a) above is also satisfied.

6. Preventing Denial-of-service Attacks

The basic 5-message protocol is susceptible to denial of service attacks. By sending a random number to a CN , an attacker can force a CN to perform a Diffie-Hellman exponentiation and an encryption operation. The CN will also create state at this point. By continuously initiating sessions with a CN , an attacker can exhaust its computation and memory resources. One way of preventing this attack would be to use “cookies”, a technique originally proposed by Karn and Simpson in [26]. Upon receiving the first message, CN replies with a cookie (which could be a keyed hash of the received exponential concatenated with a timestamp and the IP address of the sender as used in the IKE-SIGMA protocol [27]). The sender then sends the cookie back to CN proving that it is capable of receiving messages at the IP address it is claiming as its own. Thus, two additional messages are exchanged between MN and CN after the first message of the original protocol.

A useful property of our protocol is that since home agents do not create state, memory denial of service attacks are not possible on the home agents. Preventing computation denial of service attacks on the home agents reduces to the problem of detecting, without performing expensive computations, whether a message has been replayed. Replay attacks on HA_{CN} can be prevented by including a sequence number or timestamp inside the encryption in message 2 of the protocol. HA_{CN} will accept a message from CN only if the sequence number is greater than the last sequence number received from the same CN . Preventing denial of service attacks on HA_{MN} without adding extra messages, appears to be more difficult. Adding a sequence number to message 3 and including it in the signature alleviates the problem somewhat. A replayed message will be detected after performing one public key operation instead of two. Adding two extra messages for exchanging cookies, of course, solves the problem.

7. Trust Infrastructure and Authentication

The protocol in Section 4 relies on the existence of a PKI. However, it can be easily adapted to work in any infrastructure which provides a chain of trust between mobile and correspondent nodes. In particular, if a AAA infrastructure [17, 18] is deployed in which MN and CN can authenticate themselves to AAA_{MN} and AAA_{CN} respectively, and the two servers share a secret K_A , the only change required is to let the AAA servers take on the role of the home agents and replace message 3 of the original protocol by the following message:

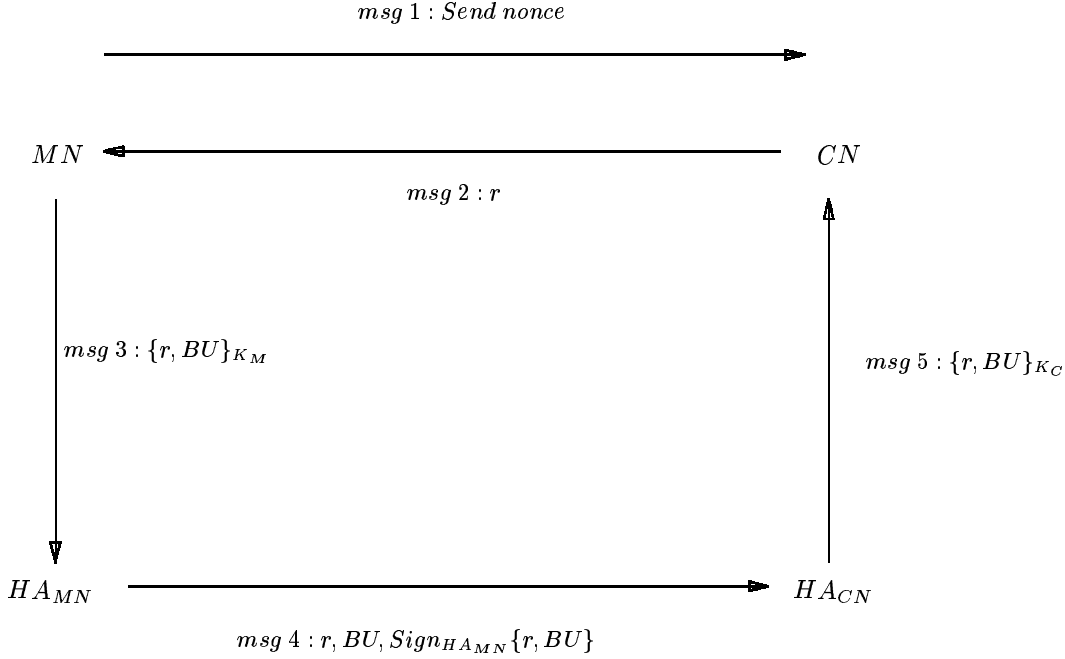


Figure 2. Protocol for Unidirectional Trust Chain

$AAA_{CN} \rightarrow AAA_{MN} : \{IP_{CN}, IP_{MN}, g^y, s_{CN}\}_{K_A}$

In the remainder of this section, we consider four alternative trust infrastructures and describe how authentication can be achieved in each one. In what follows, we assume throughout that MN and CN can authenticate themselves to $HAMN$ and $HACN$ respectively. We distinguish between the following four trust infrastructures:

1. MN and CN can directly authenticate each other.

A special case is when MN and CN share a secret key. It is then possible for MN to send an authenticated binding update to CN by executing the authentication phase protocol of Section 4.2.2. This protocol does not require the involvement of any third party and is also computationally inexpensive. Another case would be when both MN and CN possess signing-verification key pairs. However, this option is not too attractive since public key cryptography is computationally intensive and may be beyond the capabilities of mobile nodes.

2. There exists a bidirectional chain of trust between MN and CN .

An example is when $HAMN$ and $HACN$ can authenticate each other using public key cryptography. The bidirectional chain of trust can be used to set up a shared secret between MN and CN by executing the initialization phase protocol of Section 4.2.1. This protocol is more expensive than the previous one: it in-

volves home agents and requires them to perform public key operations. However, the initialization phase needs to be executed only once per $MN - CN$ pair. Once executed MN and CN end up with a shared secret and can directly authenticate each other.

3. There exists a unidirectional chain of trust from MN to CN .

Consider, for example, the case when $HAMN$ can authenticate herself to $HACN$ but not vice versa. MN can then use the trust chain $MN \rightarrow HAMN \rightarrow HACN \rightarrow CN$ to send an authenticated binding update as shown in the protocol in Figure 2. This protocol provides sender authentication and data integrity. In message 1, MN requests CN to send a nonce. The nonce, r , sent in message 2 is used to prevent replay attacks. We note, however, that this protocol places higher computation demands on the home agents. The home agents have to perform public key operations everytime a binding update is sent. It is not possible to set up a shared secret between MN and CN since CN has no way of authenticating herself to MN .

4. There is no chain of trust from MN to CN .

In the absence of a chain of trust from MN to CN , it does not appear that strong authentication is possible. However, we can design a protocol with a narrower window of vulnerability than one which provides no authentication at all. One way to achieve this

would be for two home agents to set up a shared secret the first time they need to communicate with each other. This could be done, e.g., by an unauthenticated Diffie-Hellman key exchange. Subsequently, they can authenticate each other using this key. Thus, assuming that this operation is performed securely, all future transactions between the nodes of these home agents can be secured. A similar security property could be obtained by exchanging public keys without having them certified by a Certification Authority.

We observe that there is a clear connection between infrastructural support and the computational requirements of the corresponding authentication protocol: as we relax our assumptions about the infrastructure, protocol participants have to do a greater amount of computation to complete the authentication process. The amount of computation gradually increases from case 1 to 3 in the list above. We also note that if a home agent possesses a public signing key, then all her mobile nodes can send authenticated binding updates even if the corresponding node's home agent does not possess a public key. This gives some incentive for a home agent to acquire a public key certificate. However, if the corresponding node's home agent does not possess a public key, the mobile node's home agent has to perform public key operations everytime a binding update is sent. This raises the issue of scalability. If a home agent is serving a number of mobile nodes, its computational resources might be exhausted. The trust infrastructure presented in case 2 therefore seems to be best for this application.

8. Conclusions

In response to the requirement that all location information about a mobile node in IPv6 should be authenticated [1, 3.1], we have proposed a protocol for authenticating binding updates. The other requirements are taken into account as well: the computational load on the nodes and on the routers is minimized, by eliminating expensive encryption operations and keeping the number of messages at minimum. The appropriate extensions preventing denial-of-service attacks are also suggested.

We have formally verified the correctness of the protocol using the finite-state analysis tool Mur ϕ . The fact that the Mur ϕ analysis did not uncover any bugs gives us increased confidence in the correctness of the protocol. In previous work, Mur ϕ has been used to analyze the security properties of several protocols. However, this is the first Diffie-Hellman based key exchange protocol that has been analyzed with Mur ϕ . We used a new modelling technique to capture the two most important properties of the Diffie-Hellman exchange: the computational hardness property which ensures that an intruder is not able to compute the DH

secret from the public exponentials; and the commutativity property which guarantees that both participants compute the same shared secret.

Finally, we have addressed the most difficult adoption issue for authenticated Mobile IPv6: reliance on authentication infrastructure. Besides the base protocol which assumes the existence of a bidirectional trust chain between mobile and correspondent nodes, we considered three other trust infrastructures and presented variations of the initialization phase that are appropriate to each scenario. The update phase of our protocol remains the same in each case since the basis for authenticated update is the shared secret established in the initialization phase. We believe that our protocol addresses the main issues in Mobile IPv6 authentication and makes best use of whatever infrastructure is available. In particular, the SSH-like scheme, using unauthenticated key exchange between home agents the first time these home agents participate in Mobile IP, might be the most practical one in the near term.

References

- [1] D. Johnson, C. Perkins. Mobility Support in IPv6. *Internet Draft*, July 2001.
- [2] A. Mankin, B. Patil, D. Harkins, E. Nordmark, P. Nikander, P. Roberts, T. Narten. Threat Models introduced by Mobile IPv6 and Requirements for Security in Mobile IPv6. *Internet Draft*, October 2001.
- [3] S. Bradner, A. Mankin, J.I. Schiller. A Framework for Purpose Built Keys (PBK). *Internet Draft*, February 2001.
- [4] F. Le, S.M. Faccin. Dynamic Diffie Hellman based Key Distribution for Mobile IPv6. *Internet Draft*, April 2001.
- [5] J. M. Galvin. Public Key Distribution with Secure DNS. *In Proc. 6th USENIX Unix Security Symposium*, 1996.
- [6] G. Ateniese, S. Mangard. A New Approach to DNS Security. *In Proc. 8th ACM Conference on Computer and Communications Security*, 2001.
- [7] D. Dill. The Mur ϕ Verification System. *In Proc. 8th International Conference on Computer Aided Verification*, pages 390-393, 1996.
- [8] R. Kemmerer, C. Meadows, J. Millen. Three Systems for Cryptographic Protocol Analysis. *In Journal of Cryptography*, 7(2):79-130, 1994.
- [9] A.W.Roscoe. Modelling and Verifying Key Exchange Protocols using CSP and FDR. *In Proc. 8th Computer Security Foundations Workshop*, pages 98-107, 1995.

- [10] C. Meadows. The NRL Protocol Analyzer: An Overview. *In Journal of Logic Programming*, 26(2):113-131, 1996.
- [11] D. Bolignano. Towards a mechanization of cryptographic protocol verification. *In Proc. 9th International Conference on Computer Aided Verification*, 131-142, 1997.
- [12] L. Paulson. The inductive approach to verifying cryptographic protocols. *In Journal of Computer Security*, 6:85-128, 1998. 131-142, 1997.
- [13] J.C. Mitchell, M. Mitchell, U. Stern . Automated Analysis of Cryptographic Protocols Using Mur ϕ . *In Proc. IEEE Symposium on Security and Privacy*, pages 141-153, 1997.
- [14] J.C. Mitchell, V. Shmatikov, U. Stern . Finite-State Analysis of SSL 3.0. *In Proc. 7th USENIX Security Symposium*, pages 201-216, 1998.
- [15] V. Shmatikov, J.C. Mitchell. Analysis of a Fair Exchange Protocol. *In Proc. 7th Annual Symposium on Network and Distributed System Security*, pages 119-128, 2000.
- [16] V. Shmatikov, J.C. Mitchell. Analysis of a Abuse-Free Contract Signing Protocol. *In Proc. Financial Cryptography*, 2000.
- [17] C. Rigney, A. Rubens, W. Simpson, S. Willens. Remote Authentication Dial In User Service (RADIUS). *RFC 2865*, June 2000.
- [18] P. R. Calhoun, H. Akhtar, J. Arrko, E. Guttman, A.C. Rubens, G. Zorn. Diameter Base Protocol. *Internet Draft*, November 2001.
- [19] W. Diffie, M. E. Hellman. New Directions in Cryptography. *IEEE Transactions on Information Theory*, 22(6):644-654, 1976.
- [20] A. J. Menezes, P. C. van Oorschot, S. A. Vanstone. Handbook of Applied Cryptography. *CRC Press*, 1996.
- [21] W. Diffie, P. C. Van Oorschot, M. J. Wiener. Authentication and Authenticated Key Exchanges. *Designs, Codes and Cryptography*, 2:107-125, 1992.
- [22] H. Krawczyk, M. Bellare, R. Canetti. HMAC: Keyed-Hashing for Message Authentication. *RFC 2104*, February 1997.
- [23] D. Dill, S. Park, A. G. Nowatzyk. Formal Specification of Abstract Memory Models. *In Symposium on Research on Integrated Systems*, pages 38-52, 1993.
- [24] U. Stern, D. Dill. Automatic Verification of the SCI Cache Coherence Protocol. *In Advanced Research Working Conference on Correct Hardware Design and Verification Methods*, pages 21-34, 1995.
- [25] V. Shmatikov, U. Stern . Efficient Finite-State Analysis for Large Security Protocols. *In Proc. 11th IEEE Computer Security Foundations Workshop*, pages 106-115, 1998.
- [26] P. Karn, W. Simpson. Photuris: Extended Schemes and Attributes. *RFC 2523*, March 1999.
- [27] H. Krawczyk. The IKE-SIGMA Protocol. *Internet Draft*, November 2001.