

Automated Design of Security Protocols

CHEN HAO, JOHN A CLARK AND JEREMY L JACOB ^a

^aDepartment of Computer Science, University of York, York, YO10 5DD, UK. email: {chenhao,jac,jeremy}@cs.york.ac.uk

Security protocols play an important role in modern communications. However, security protocol development is a delicate task, and experience shows that computer security protocols are notoriously difficult to get right. Recently, Clark and Jacob provided a framework for automatic protocol generation based on combinatorial optimization techniques and the symmetric key part of BAN logic. This paper shows how such an approach can be further developed to encompass the full BAN logic without loss of efficiency and thereby synthesize public key protocols and hybrid protocols.

1 Introduction

When we look into published security protocols, we find that many of these protocols do not succeed in their stated or implied goals. Although inappropriate use of cryptography may pose problems, it is clear that many protocols suffer problems which have nothing to do with the strength of the cryptography used. Many of the errors arise from the inappropriate structure of the message exchange. As a result, many existing protocols are susceptible to various kinds of attacks, which are independent of the weaknesses of the cryptosystem employed.

Various formalisms and tools have been brought to bear on the problem. However, it would seem that automated support in this area is largely limited to the analysis and formal verification of existing protocols; there is little work in automatic protocols synthesis. The first work on automatic protocols synthesis using metaheuristic search was presented by Clark and Jacob [6, 7]. However, its application was limited in various ways. Most obviously, principals were allowed to use only symmetric key encryption. In this paper, we show how we have extended this technique to allow public key encryption and hybrid schemes. We also report the results of extensive experimentation with the technique.

2 Protocols and Belief Logic

Security protocols are designed to let principals communicate securely over an insecure network. Security requirements include:

1. **Secrecy.** An intruder should not be able to read the contents of messages intended for others.
2. **Authenticity.** If a message appears to be from a principal A for an identified purpose, then A sent that message for that purpose.
3. **Non-repudiation.** If A sent a message, that principal cannot later deny it.

Anderson and Needham show that security protocol development is a delicate task, and computer security protocols are notoriously difficult to get right [2]. Recent approaches to use formal methods in the design of security protocols include finite-state model checking and belief logic. In this paper, we concentrate on belief logics, which formalize what a principal may infer from messages received, taking as our example the first such logic, BAN [3].

In 1989, Burrows, Abadi and Needham developed a belief logic (BAN logic) that could be used to reason about protocol security. Although their work has aroused much debate, the BAN logic is a milestone in the area of security protocol design and analysis. BAN logic focuses on the beliefs of honest parties involved in the protocols and on the evolution of these beliefs as a consequence of communication. The original BAN logic allows short, abstract proofs. It has identified some protocol flaws but missed others [5]. As a result, a number of variations and enhancements of the BAN logic have been developed. These new belief logics, such as GNY logic [8] and SVO logic [15], address some weaknesses of BAN logic but sacrifice its simplicity. The work described here is based on BAN, but we believe it could easily be ported to other logics. Below is a brief introduction to the notation and inference rules used in BAN logic.

2.1 Notions and Notations

Idealized Protocols In much of the literature, security protocols have not been expressed in a formal manner. Such descriptions must be converted to formal descriptions if formal analysis is to take place. In BAN logic literature the abstractions that are analyzed are termed *idealized* protocols. In concrete protocols, principals maintain data items (e.g. keys) and communicate some of these items using messages with an agreed format. On receiving a message the receiver will update its state in some agreed way and carry out other agreed actions. With idealized protocols, principals maintain and communicate *beliefs*. Thus, rather than holding a key K_{ab} for session communication between A and B and distributing that key to A and B , a server S would hold the *belief* that the key K_{ab} was good for communicating between A and B (denoted $A \xleftrightarrow{K_{ab}} B$) and include that belief in messages to A and B . The logic indicates how a receiver should update its belief states on receipt of a message.

Encryption and Keys All messages in BAN logic are encrypted. Unencrypted messages sent over an insecure network provide no guarantees of any kind, because an intruder may easily alter cleartext. In practice, unencrypted concrete messages may be used as signals to cause encrypted messages to be sent, but they do not contribute to principals' beliefs. Since we work at the abstract BAN level some of our protocols have principals sending messages apparently without stimulus. Supplying such stimulus is a concrete *implementation* issue.

Nonce All beliefs held in the current run of a protocol are stable for the entirety of the protocol; however, beliefs held in the past are not necessarily carried forward into the present. Therefore, it is important for principals involved in a protocol to determine that messages they receive really have been created as part of the current run of the protocol. This is typically achieved by the inclusion in messages of data to bind messages to the current run. This data takes the form of numbers generated to be used only once (for bindings to the current run). These numbers used only once are commonly called *nonces*. If a principal generates a nonce for the current protocol run and receives messages that contain it, this principal may deduce that these messages have been created after the nonce was generated. An alternative to nonces are timestamps, which can also make the receiver believe the messages have been generated recently.

Basic Notation The language of BAN consists of the following expressions:

Believes. The assertion $P \models X$ means P believes the formula X . P may act as if X is true.

Sees. The assertion $P \triangleleft X$ means P sees X . Someone has sent a message containing X to P , and P can read and repeat X ; this may require decryption.

Once Said. The assertion $P \sim X$ means P once said X . The principal P at some time sent a message including the statement X . It is assumed that P believed X when it sent the message.

Jurisdiction. The assertion $P \models X$ means P has jurisdiction over X . The principal P is an authority on X and should be trusted on this matter. An example of jurisdiction is that principals may believe that a key distribution server has jurisdiction over statements about the quality of keys.

Fresh. The assertion $\#(X)$ means the formula X is fresh, that is to say, X has not been sent in a message at any time before the current run of the protocol. This is usually true for nonces.

Key Goodness. The assertion $P \xleftrightarrow{K} Q$ means K is a good key for communication between P and Q . That is, the key K has not been revealed to any principal other than P or Q .

Public key. The assertion $\overset{K}{\mapsto} P$ stands for the principal P having a public key K . The matching secret key (denoted by K^{-1}) will never be revealed to any principal other than P .

Secret. The assertion $P \overset{X}{\rightleftharpoons} Q$ means the formula X is a secret known only to P and Q , and possibly to principals trusted by them. Only P and Q may use X to prove their identities to one another. An example of a shared secret is a password.

Encryption. The assertion $\{X\}_K$ means the formula X encrypted under the key K . Principals can recognize their own messages. Encrypted messages are uniquely readable and verifiable as such by holders of the right keys. Similarly, encrypted messages can be created only by a principal with the appropriate keys.

Combined. The assertion $\langle X \rangle_Y$ represents X combined with the formula Y ; it is intended that Y be a secret, and that its presence prove the identity of whoever utters $\langle X \rangle_Y$.

2.2 Inference Rules

When a principal receives a message, the logic provides inference rules that indicate what new beliefs this principal may infer from the message contents. The major inference rules are given below.

Message Meaning Rules The message meaning rules explain how to derive beliefs about the origin of messages. Two of the three concern the interpretation of encrypted messages, and the third concerns the interpretation of messages with secrets.

$$\frac{P \equiv P \xleftrightarrow{K} Q, P \triangleleft \{X\}_K}{P \equiv Q \mid \sim X}$$

That is, if principal P believes the key K is shared only with principal Q , and sees a message X encrypted under that key K , then P may conclude that this message X was created by Q , who ‘once said’ its contents X .¹

Similarly, for public keys:

$$\frac{P \equiv P \xrightarrow{K} Q, P \triangleleft \{X\}_{K^{-1}}}{P \equiv Q \mid \sim X}$$

That is, if principal P believes that the key K is Q ’s public key and it receives a message $\{X\}_{K^{-1}}$ encrypted under Q ’s corresponding (private) inverse key K^{-1} , then P may conclude that principal Q once said the contents of the message.

For shared secrets:

$$\frac{P \equiv P \xleftrightarrow{Y} Q, P \triangleleft \langle X \rangle_Y}{P \equiv Q \mid \sim X}$$

That is, if principal P believes that the secret Y is shared only with Q and sees $\langle X \rangle_Y$, then P believes that Q once said X .

Nonce Verification Rule The nonce verification rule expresses how a principal’s view of a message changes when it determines that the message is part of the current protocol run.

$$\frac{P \equiv \#(X), P \equiv Q \mid \sim X}{P \equiv Q \mid \equiv X}$$

That is, if P believes that X is fresh and that Q once said X , then P believes that Q has said X during the current run of protocol, and hence that Q believes X at present. In order to apply this rule, X should not contain any encrypted text. The nonce verification rule is the only way of ‘promoting’ once said assertion to actual belief.

Jurisdiction Rule The jurisdiction rule captures the notion that some principals are trusted to carry out certain tasks and make particular judgments.

$$\frac{P \equiv Q \mid \equiv X, P \equiv Q \mid \Rightarrow X}{P \equiv X}$$

That is, if principal P believes that Q believes X , and also believes that Q has jurisdiction over X , then P should believe X too.

In this paper, we also need some smaller rules such as the ability to deduce $A \mid \equiv \#(X, Y)$ from $A \mid \equiv \#(X)$. But we shall omit these here. Further details of these inference rules can be found in Burrows, Abadi and Needham’s paper [3].

¹In BAN logic, it is assumed that principals can recognize messages they themselves have created and take appropriate action when they receive such messages. We shall interpret $P \triangleleft \{X\}_K$ as ‘ P sees message $\{X\}_K$ and P knows that it did not create $\{X\}_K$ ’ itself.

Initial Assumptions	
$S \models \overset{K_a}{\mapsto} A$	S believes that K_a really is A 's public key
$S \models \overset{K_b}{\mapsto} B$	S believes that K_b really is B 's public key
$S \models \overset{K_s^{-1}}{\mapsto} S$	S believes that K_s^{-1} is its own private key
$A \models \overset{K_s}{\mapsto} S$	A believes that K_s is the public key of S
$A \models S \Rightarrow \overset{K_b}{\mapsto} B$	A trusts S to provide B 's public key, that is, A believes that S has jurisdiction over B 's public key
$A \models N_a$	A believes that a particular number N_a is a well-formed nonce.
$A \models \#(N_a)$	A believes that nonce N_a is actually fresh.
$A \models \overset{K_a^{-1}}{\mapsto} A$	A believes that K_a^{-1} is its own private key.

Protocol Goal	
$A \models \overset{K_b}{\mapsto} B$	A believes that K_b is B 's public key

Protocol	
1. $A \rightarrow S : \{N_a\}_{K_a^{-1}}$	
2. $S \rightarrow A : \left\{ A \mid \sim N_a, \overset{K_b}{\mapsto} B \right\}_{K_s^{-1}}$	

Figure 1: Initial assumptions, a goal and a feasible protocol

2.3 Illustrative Example

Figure 1 gives a set of initial assumptions held by principal A and a key distribution server S , and a feasible protocol.

A believes N_a is a well formed nonce and may include it in the first message. This message is encrypted with its private key K_a^{-1} . When the server S sees (receives) this encrypted message, it can use A 's public key to decrypt it and deduce $A \mid \sim N_a$, that is A once said N_a , via the *Message Meaning Rule*.

Now, S may reply to A with the second message that contains two of its current beliefs: the newly derived belief A once said N_a and an initial assumption $\overset{K_b}{\mapsto} B$. S encrypts this message using its private key K_s^{-1} . Once A sees this message, he may decrypt it to reveal its contents. Using the *Message Meaning Rule*, A concludes $S \mid \sim \overset{K_b}{\mapsto} B$, that is S once said K_b is B 's public key. In the meantime, A may also conclude $S \mid \sim A \mid \sim N_a$, that is, S once said that A once said N_a . This message contains an assertion involving N_a , a nonce A believes to be fresh, so A may conclude the whole message is a fresh one. Then A may deduce that S believes the whole message using the *Nonce Verification Rule*. In detail, A concludes $S \models A \mid \sim N_a$ and also $S \models \overset{K_b}{\mapsto} B$. Since A believes that S has jurisdiction over B 's public key, A may now believe $\overset{K_b}{\mapsto} B$ using the *Jurisdiction Rule*.

3 Search Strategy – Metaheuristic Techniques

What we wish to do, given some assumptions and goals, is to find protocols that achieve the goals from the assumptions. That is, we wish to search the space of feasible protocols for ones satisfying a specification. Any series of honest exchanges between two or more principals defines a feasible (with respect to the logic) protocol. This is the set of feasible protocols that we consider as the design space. It is clear that this space grows exponentially as the number of messages or the number of principals rise. The choices of the belief contents of messages introduce further combinatorial complexity. For a technique to be scalable it cannot be based on simple enumeration.

Metaheuristics are widely used to solve important practical combinatorial optimization problems [13]. The role of metaheuristic search is to exchange guarantees of optimality for computational tractability.

Examples of metaheuristics include simulated annealing (SA) [10], tabu search (TS) [1], genetic algorithms (GA) [9], and ant colony optimization (ACO) [4]. The results reported here were obtained using simulated annealing; genetic algorithms are currently under investigation.

To use a search approach for protocol synthesis, we need to provide:

- a characterization of the design space – how protocols are represented and how valid and invalid protocols can be distinguished.
- a fitness function – if we wish to obtain the ‘best’ or just ‘good’ protocols, we must characterize precisely how ‘good’ a candidate protocol is.
- a search strategy – when exhaustive search is impractical, we need to provide a strategy for searching the design space that can locate good protocols within reasonable computational time.

We discussed the first issue in section 2. In this section, we show details of how to implement valid protocols within an optimization framework and the fitness functions that we use to guide the search to a solution.

3.1 Protocol Representation and a Move Function

A BAN protocol is represented in our search algorithm as a sequence of M messages, each of which is represented by a non-negative integer sequence. A message is sent by one principal and received by another. N principals, indexed $0 \dots N - 1$, participate in the protocol. Associated with each of the principals is a vector of its current beliefs. Each of the M protocol messages is represented by $B + 3$ integers, $vs, vr, vk, vb_1, \dots, vb_B$. These represent the sender, the receiver, the key that the sender used to encrypt this message, and a series of B indices that reference beliefs currently held by the sending principal. So, the sender is $vs \bmod N$; receiver is $vr \bmod N$; key is $vk \bmod (2N + N(N - 1)/2)$ (N principals may have N private keys, N public keys, and share $N(N - 1)/2$ symmetric keys); the first belief in the message is belief $vb_1 \bmod T$ etc., where the sender has T current beliefs, indexed $0 \dots T - 1$. $belief[0]$ is the null belief (which allows us to model easily messages with fewer than B ‘real’ beliefs). The vector of the receiver’s current beliefs is updated after each message is sent (see below). In this way, an arbitrary sequence of integers can be interpreted as a feasible protocol (by construction, senders only ever send beliefs they actually hold). This allows a very simple move strategy for local search – simply randomly perturb any of the integers involved in any message. Although the interpreted protocol may be feasible, it may not satisfy our required goals. The fitness function below measures how close it comes to achieving the required goals and our search seeks to find a protocol that satisfies all these goals.

3.2 Executing a Protocol

Assume a protocol consists of M messages, each of which consists of B beliefs, and we start from the very beginning of this protocol. Firstly, we should initialize the belief state of each principal involved in this protocol with its initial assumptions. Then, for each message in this protocol, we follow the steps below.

1. Determine the sender, receiver, and the key under which the current message is encrypted. If this key is an appropriate one for communication between the sender and the receiver, then proceed with the rest of the current message, else ignore this message and proceed to the next message. The method of decoding the sender, receiver, and key is indicated in section 3.1
2. Decode each of the B beliefs corresponding to the current message. For instance, the first belief in the message is $vb_1 \bmod T$, where the sender currently holds T sendable beliefs.
3. Update the receiver’s beliefs vector by applying the message meaning rule, the nonce verification rule, and the jurisdiction rule in that order. Here we demonstrate what a principal P will do after it receives a message $\{X, P \sim N_p\}$ from another principal Q . Firstly, $Q \sim X$ and $Q \sim P \sim N_p$ are added to P ’s belief vector (this represents $P \equiv Q \sim X$ and $P \equiv Q \sim P \sim N_p$). This, together with (1) above implements the message meaning rule. After this, P examines the set of received beliefs to see whether any of the beliefs contain a component that it believes to be fresh. In this case, P receives the belief $P \sim N_p$, and if P believes the nonce N_p is fresh, then the whole message is

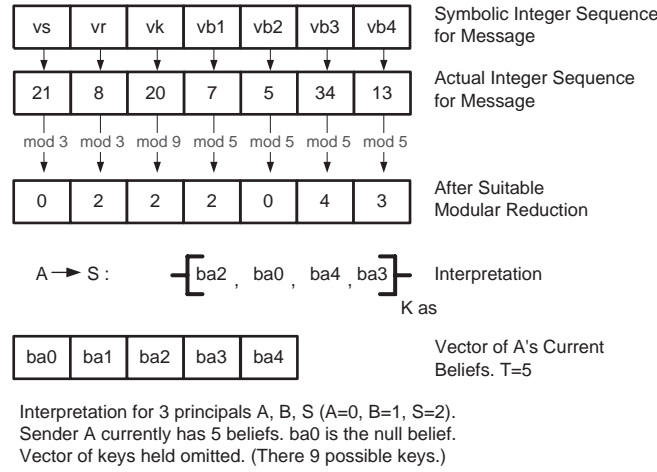


Figure 2: Interpreting an Integer sequence

regarded as fresh. If the message is fresh then the nonce verification rule is applied to add $Q \mid\equiv X$ and $Q \mid\equiv P \sim N_p$. Similarly, the jurisdiction rule now may be applied to deduce further beliefs until no further beliefs can be created.

4. Record the number of required goals achieved after this message has been analyzed.

Once a protocol has been executed in the above way, the fitness of this protocol can be calculated as given in the following section.

3.3 The Fitness Function

The fitness function is used to guide the search for a ‘good’ solution, that is, the fitness function must tell how ‘good’ a candidate solution is. We use fitness functions for a protocol of the form:

$$\sum_{i=1}^M w_i * g_i$$

The w_i are weightings and g_i is the number of required goals achieved after message i . In this paper, we use several weighting strategies for setting the weights w_i that are detailed in table 1. These weighting strategies were first used by Clark and Jacob. Here we only introduce these weighting strategies briefly, further details of these strategies can be found in [6, 7]. Note that the fitness function used rewards cumulatively.² If a goal becomes satisfied after some message it is also satisfied after all subsequent messages. Thus, the g_i form a monotone non-decreasing sequence. The user specifies the number M of messages in a protocol.

- *Early Credit (EC)*. The weights are monotonically decreasing with i . The notion is that satisfying goals early should be rewarded.
- *Uniform Credit (UC)*. All the weights are the same.
- *Delayed Gratification (DG)*. The weights are monotonically increasing. This captures the idea that early satisfaction of goals may not necessarily be a good thing.
- *Advanced Delayed Gratification (ADG)*. The weights are monotonically increasing and no credit is given immediately for satisfying goals in the initial exchanges.
- *Uniform Delayed Gratification (UDG)*. No credit is given immediately for satisfying goals in the initial exchanges and later weights are equal and positive.

²Others approaches are clearly possible.

Weight	Strategy					
	EC	UC	DG	ADG	UDG	DJ
w_1	2000	500	50	0	0	0
w_2	1000	500	100	0	0	0
w_3	500	500	200	200	1000	0
w_4	200	500	500	500	1000	0
w_5	100	500	1000	1000	1000	0
w_6	50	500	2000	2000	1000	0
w_7	25	500	4000	4000	1000	1000

Table 1: Weighting Strategies

- *Destination Judgment (DJ)*. Only the final weights are non-zero. It does not matter how you satisfy goals, the important thing is how many you satisfy in the end.

3.4 Optimization Techniques

In this paper we have used the well-established technique of simulated annealing [10], though our implementation allows rapid interchange of optimization techniques. The annealing approach is the standard one with a geometric cooling rate of 0.97. The number of attempted moves at each temperature is 400, with a maximum of 1000 iterations (temperature reductions) and maximum number of 50 consecutive unproductive iterations (i.e. with no move being accepted). In the interests of brevity we assume the audience is familiar with the standard annealing algorithm.

4 Experimental Method and Results

The original Clark-Jacob technique and supporting tools dealt only with symmetric key protocols. The technique and tools have now been extended to allow symmetric, public and hybrid protocols to be synthesised. This section reports the results of applying the extended technique described above to the derivation of three-party key distribution protocols. The results consist of two aspects, the protocols and the success fractions. This section is organized by different sorts of protocols.

4.1 Public key Protocols

Initial Assumptions Three principals involved in this key distribution protocol are A , B and S . As a key distribution server, S holds all the other principals' public keys and its own private key. A and B both have the server's public key and their own private keys. They also maintain their own nonces that they believe to be fresh. A and B each believes that S is to be trusted on the other's public key. All the assumptions are listed below.

$$\begin{aligned}
S &| \equiv \overset{K_a}{\mapsto} A, S | \equiv \overset{K_b}{\mapsto} B, S | \equiv \overset{K_s^{-1}}{\mapsto} S; \\
A &| \equiv \overset{K_s}{\mapsto} S, A | \equiv \overset{K_a^{-1}}{\mapsto} A, A | \equiv N_a, A | \equiv \#(N_a), A | \equiv S \Rightarrow \overset{K_b}{\mapsto} B; \\
B &| \equiv \overset{K_s}{\mapsto} S, B | \equiv \overset{K_b^{-1}}{\mapsto} B, B | \equiv N_b, B | \equiv \#(N_b), B | \equiv S \Rightarrow \overset{K_a}{\mapsto} A.
\end{aligned}$$

Goals At the end of the protocol run, both A and B must believe that they hold each other's public key. The other two goals require that each of them believes the other believes its public key is good.

$$\begin{aligned}
A &| \equiv \overset{K_b}{\mapsto} B, B | \equiv \overset{K_a}{\mapsto} A; \\
A &| \equiv B | \equiv \overset{K_a}{\mapsto} A, B | \equiv A | \equiv \overset{K_b}{\mapsto} B.
\end{aligned}$$

Results and Statistics Twenty runs of the program were carried out for each fitness function strategy. In our program, the annealing parameters given in section 3.4 were used. Figure 3 shows one of the public key protocols generated by the program.

In the rest of this paper, only the core security relevant components of a protocol are presented. That is, our descriptions of protocols do not include those belief components that do not contribute to the predefined goals. In addition, redundant beliefs (where the same beliefs are included twice or more in one message) have also been removed. Currently, these ‘junk’ beliefs are removed by hand; automating their removal is under investigation.

1. $A \rightarrow S : \{N_a\}_{K_a^{-1}}$
2. $S \rightarrow A : \left\{ A | \sim N_a, \overset{K_b}{\mapsto} B \right\}_{K_s^{-1}}$
3. $B \rightarrow S : \{N_b\}_{K_b^{-1}}$
4. $S \rightarrow B : \left\{ B | \sim N_b, \overset{K_a}{\mapsto} A \right\}_{K_s^{-1}}$
5. $B \rightarrow A : \{N_b\}_{K_b^{-1}}$
6. $A \rightarrow B : \left\{ B | \sim N_b, N_a, \overset{K_b}{\mapsto} B \right\}_{K_a^{-1}}$
7. $B \rightarrow A : \left\{ A | \sim N_a, \overset{K_a}{\mapsto} A \right\}_{K_b^{-1}}$

Figure 3: A public key protocol generated during experimentation

The search rapidly established the first two beliefs $A \equiv \overset{K_b}{\mapsto} B$ and $B \equiv \overset{K_a}{\mapsto} A$ after 4 messages. In order to achieve the third goal $B \equiv A \equiv \overset{K_b}{\mapsto} B$, A must have knowledge of the nonce N_b and show it to B (A sends B a message including $B | \sim N_b$). In our program, sendable beliefs are either simple (e.g. N_a , $\overset{K_a}{\mapsto} A$ are both sendable) or else involve only one operator (e.g. $A | \sim N_a$, $S \mid \Rightarrow \overset{K_a}{\mapsto} A$ and so on). The only way A can acquire knowledge of N_b is to receive it in a message from B directly. (It would be possible to obtain knowledge of N_b via S , but this would require additional assumptions by A . For example, if S sent a message including $B | \sim N_b$ and $A \equiv S \mid \Rightarrow B | \sim N_b$, then the message would be considered fresh by A).

When we input $A \equiv S \mid \Rightarrow B | \sim N_b$ and $B \equiv S \mid \Rightarrow A | \sim N_a$ as two initial assumptions, our program generates protocols that achieve our goals (same as before) within 6 messages. Figure 4 shows one of these protocols.

1. $A \rightarrow S : \{N_a\}_{K_a^{-1}}$
2. $B \rightarrow S : \{N_b\}_{K_b^{-1}}$
3. $S \rightarrow A : \left\{ A | \sim N_a, B | \sim N_b, \overset{K_b}{\mapsto} B \right\}_{K_s^{-1}}$
4. $S \rightarrow B : \left\{ B | \sim N_b, A | \sim N_a, \overset{K_a}{\mapsto} A \right\}_{K_s^{-1}}$
5. $B \rightarrow A : \left\{ A | \sim N_a, \overset{K_a}{\mapsto} A \right\}_{K_b^{-1}}$
6. $A \rightarrow B : \left\{ B | \sim N_b, \overset{K_b}{\mapsto} B \right\}_{K_a^{-1}}$

Figure 4: A public key protocol with additional assumptions

From an abstract logic point of view, the more beliefs included in messages the more information the receiving principal obtains. Thus, messages with more information create a greater probability of achieving goals. We have also repeated the above experiments allowing three beliefs per message. Figure 5 gives the success fractions when four and three beliefs per message are used. As we have seen, for the original problem, in all cases except for the destination judgment (DJ), the success fractions are decreased dramatically. Another conclusion we can draw from these figures is that appropriate redundancy is actually very useful to the optimization approach. Moreover, DJ is clearly awful in both cases. Some degree of reward for early achievement is clearly useful.

Strategy	Success Fraction Four Beliefs Per Message	Success Fraction Three Beliefs Per Message
<i>EC</i>	0.70	0.40
<i>UC</i>	0.80	0.60
<i>DG</i>	0.90	0.60
<i>ADG</i>	0.85	0.60
<i>UDG</i>	0.80	0.40
<i>DJ</i>	0.05	0.05

Figure 5: Results for public key protocols with four beliefs and three beliefs per message

4.2 Public Key Protocols using Timestamps

Initial Assumptions Again, three principals involved in this sort of key distribution protocol are A , B and S , where S is a trustworthy key distribution server. Here we allow the notion of timestamps. T is, effectively, a form of nonce shared by all parties prior to the run of the protocol. The difference is that this sort of protocol relies heavily on synchronized clocks, since each principal believes that a timestamp generated elsewhere is ‘fresh’ if it has a value within a window of the receiver’s local time.

$$\begin{aligned}
A &| \equiv \xrightarrow{K_s} S, A | \equiv \xrightarrow{K_a^{-1}} A, A | \equiv T, A | \equiv \#(T), A | \equiv S | \Rightarrow \xrightarrow{K_b} B; \\
B &| \equiv \xrightarrow{K_s} S, B | \equiv \xrightarrow{K_b^{-1}} B, B | \equiv T, B | \equiv \#(T), B | \equiv S | \Rightarrow \xrightarrow{K_a} A; \\
S &| \equiv \xrightarrow{K_a} A, S | \equiv \xrightarrow{K_b} B, S | \equiv \xrightarrow{K_s^{-1}} S, S | \equiv T, S | \equiv \#(T).
\end{aligned}$$

Goals We hope this sort of protocol can achieve the following four goals (as before).

$$\begin{aligned}
A &| \equiv \xrightarrow{K_b} B, B | \equiv \xrightarrow{K_a} A; \\
A &| \equiv B | \equiv \xrightarrow{K_a} A, B | \equiv A | \equiv \xrightarrow{K_b} B.
\end{aligned}$$

Results and Statistics We use the same annealing parameters as those used in section 4.1. Figure 6 shows one of the protocols generated by our program, and Figure 7 shows the success fraction for each search strategy when we allow four beliefs in each message. As might be expected the evolved protocols have fewer messages.

1. $S \rightarrow A : \left\{ T, \xrightarrow{K_b} B \right\}_{K_s^{-1}}$
2. $S \rightarrow B : \left\{ T, \xrightarrow{K_a} A \right\}_{K_s^{-1}}$
3. $B \rightarrow A : \left\{ T, \xrightarrow{K_a} A \right\}_{K_b^{-1}}$
4. $A \rightarrow B : \left\{ T, \xrightarrow{K_b} B \right\}_{K_a^{-1}}$

Figure 6: A public key protocol using timestamps, four beliefs per message

4.3 Hybrid Protocols

We now attempt to evolve a protocol allowing the use of both symmetric and public key encryption.

Initial Assumptions Essentially, we aim to distribute a secret key using public key means. In this sort of protocol, we assume that both principals A and B can communicate with the server S via a public key. The server S will distribute a symmetric session key that will be used in further communications between A and B .

Strategy	Success Fraction
EC	0.95
UC	1.00
DG	0.90
ADG	0.65
UDG	0.85
DJ	0.60

Figure 7: Results for public key protocol using timestamps

$$\begin{aligned}
A &| \equiv \xrightarrow{K_s} S, A | \equiv \xrightarrow{K_a^{-1}} A, A | \equiv N_a, A | \equiv \#(N_a), A | \equiv S \Rightarrow A \xleftrightarrow{K_{ab}} B; \\
B &| \equiv \xrightarrow{K_s} S, B | \equiv \xrightarrow{K_b^{-1}} B, B | \equiv N_b, B | \equiv \#(N_b), B | \equiv S \Rightarrow A \xleftrightarrow{K_{ab}} B; \\
S &| \equiv \xrightarrow{K_a} A, S | \equiv \xrightarrow{K_b} B, S | \equiv \xrightarrow{K_s^{-1}} S, S | \equiv A \xleftrightarrow{K_{ab}} B.
\end{aligned}$$

Goals The four desired goals are:

$$\begin{aligned}
A &| \equiv A \xleftrightarrow{K_{ab}} B, B | \equiv A \xleftrightarrow{K_{ab}} B; \\
A &| \equiv B | \equiv A \xleftrightarrow{K_{ab}} B, B | \equiv A | \equiv A \xleftrightarrow{K_{ab}} B.
\end{aligned}$$

Results and Statistics Figure 8 shows one of the hybrid protocols generated by the program. Figure 9 shows the success fraction for each search strategy when we allow four beliefs in each message. When con-

$$\begin{aligned}
1. & A \rightarrow S : \left\{ \{N_a\}_{K_a^{-1}} \right\}_{K_s} \\
2. & S \rightarrow A : \left\{ \left\{ A | \sim N_a, A \xleftrightarrow{K_{ab}} B \right\}_{K_s^{-1}} \right\}_{K_a} \\
3. & B \rightarrow S : \left\{ \{N_b\}_{K_b^{-1}} \right\}_{K_s} \\
4. & S \rightarrow B : \left\{ \left\{ B | \sim N_b, A \xleftrightarrow{K_{ab}} B \right\}_{K_s^{-1}} \right\}_{K_b} \\
5. & B \rightarrow A : \{N_b\}_{K_{ab}} \\
6. & A \rightarrow B : \left\{ B | \sim N_b, N_a, A \xleftrightarrow{K_{ab}} B \right\}_{K_{ab}} \\
7. & B \rightarrow A : \left\{ A | \sim N_a, A \xleftrightarrow{K_{ab}} B \right\}_{K_{ab}}
\end{aligned}$$

Figure 8: A hybrid encryption protocol

sidering protocols that distribute secret information, we assume that any signed messages are then encrypted with the public key of the receiver as shown in figure 8. This is a sound assumption that is recommended by Anderson and Needham [2]. However, we can see that the nonce need not be kept secret and so the outer encryptions (using K_s) in message 1 and 3 may be removed. We can see that (default) other encryptions in message 2 and 4 are necessary (to maintain confidentiality of the shared key K_{ab}).

4.4 Extensive Experimentation

We have applied our approach to the on-line repository of security protocols “Security Protocols Open Repository” at <http://www.lsv.ens-cachan.fr/spore/>, which contains 45 protocols. Our program successfully synthesized 38 BAN protocols when given the same initial assumptions and goals as they are in the original ones. Obviously, all these 38 protocols are correct according to the BAN logic, and the successful synthesis process itself is a proof. The remaining 7 sets of assumptions and goals contain features outside of the BAN logic, and so our tool cannot be used. There are several “variations on a theme” in the library. Often, the assumptions and goals of these variants are the same. Thus, there may be

Strategy	Success Fraction
<i>EC</i>	0.80
<i>UC</i>	0.85
<i>DG</i>	0.85
<i>ADG</i>	0.60
<i>UDG</i>	0.80
<i>DJ</i>	0.05

Figure 9: Hybrid Protocols

one abstract specification and several concrete implementations. The 38 evolved concrete correspond to 23 distinct abstract specifications.

Experimentation highlighted difficulties with repeated authentication. In a typical repeated authentication protocol, the first (preliminary) part typically distributes a ‘ticket’ to a principal. This ticket can be used by the principal to achieve some authentication task. This may be described as the full initial run of the protocol. However, the ticket will typically be used several more times (within a specified lifetime). A well-known repeated authentication protocol is the Neumann Stubblebine protocol given in Appendix A.17.

The logic (and so our tools) has difficulties with the use of repeated ‘tickets’. We have simply evolved protocols to meet the goals of the first authentication run. It is generally possible to address the repeated parts of the protocol, provided each presentation of a ticket is regarded as ‘fresh enough’ (or simply ‘fresh’) if its lifetime has not expired. With such an approach, the evolution of mutual authenticating nonce exchanges is generally trivial. Nevertheless, our tools currently do not allow keys of the form K_{pp} (a key known only to P and used to seal tickets to be presented repeatedly to P).

5 Conclusions and Further Work

The above work shows that the original Clark-Jacob approach for the symmetric case can be successfully extended to allow public key and hybrid cryptographic schemes. The protocols generated, although simple, are typical abstractions of protocols in the literature. The ease with which the approach generated protocols satisfying realistic goals merits further investigation of the technique. Experimentation and our general knowledge of protocol verification techniques have allowed us to identify numerous possible improvements to the approach and tool support. These are outlined below.

A more sophisticated logic (SVO seems a promising candidate) should be adopted to increase design choice and give greater confidence in the practical security of evolved protocols. As far as actual freedom from security flaws is concerned, we are very much at the mercy of the logic we choose. We have expanded the previously used subset of BAN logic to allow public key and hybrid protocols to be evolved. We need now to address weaknesses in the BAN logic itself. Similarly, allowing more sophisticated beliefs to be communicated in messages should allow a wider range of protocols to be evolved.

Non-functional properties such as efficiency are an important consideration for most security protocol designers and should be incorporated into our design synthesis approach. Efficiency has not been ignored completely in the current approach — the cumulative reward nature of the fitness function generally favors shorter protocols — but this is somewhat indirect and does not address crucial issues such as amount of encryption etc. Automatic refinement to a more detailed representation (code, for example) would be a significant enhancement and would greatly facilitate inclusion of non-functional issues.

We may need to exploit the approach taken to the fullest extent possible. The model checking approaches [11, 12] are distinctly limited, e.g. three or four messages, in the size of the protocols they can produce. In this paper we have presented seven-message protocols and some nine-message protocols were demonstrated by Clark and Jacob[7]. We currently do not know the limits of the optimization approaches. Earlier work [7] showed that the protocol generation problem suffers from combinatorial explosion. As the complexity of the underlying logic increases, so does the magnitude of the search problem. Simulated annealing and genetic algorithms may not be the best optimization techniques to use. Others should be considered.

The work here shows that metaheuristic search approaches to secure protocol synthesis are potentially powerful and have the benefit of being rapid. Our tools could generate candidate protocols rapidly and

concrete refinements of them could be subjected to more detailed and sophisticated analysis (such as that provided by current model checking approaches). This would provide an interesting synthesis of current techniques.

Our experiments have tested the approach's ability to generate protocols for existing and fairly standard requirements. It seems suited to such tasks. It will be interesting to see whether metaheuristic approaches will be able to produce protocols for novel or highly complex requirements. Will there be any surprises?

References

- [1] D. de Werra A. Hertz, E. Taillard. A tutorial on tabu search. In *Proc. of Giornate di Lavoro AIRO'95 (Enterprise Systems: Management of Technological and Organizational Changes)*, pages 13–24, Italy, 1995.
- [2] Martín Abadi and Roger Needham. Prudent engineering practice for cryptographic protocols. *IEEE Transactions on Software Engineering*, 22(1):6–15, 1996.
- [3] Michael Burrows, Martín Abadi, and Roger Needham. A logic of authentication. Technical Report 39, Digital Systems Research Center, February 1989.
- [4] Eric Bonabeau, Marco Dorigo, and Guy Theraulaz. *Swarm Intelligence: From Natural to Artificial Systems (Santa Fe Institute Studies on the Sciences of Complexity)*. Oxford University Press Inc, USA; ISBN: 0195131592, 1999. ISBN: 0195131592.
- [5] Colin Boyd and Wenbo Mao. On a limitation of BAN logic. In *Advances in Cryptology - EUROCRYPT '93: Workshop on the Theory and Application of Cryptographic Techniques, Lecture Notes in Computer Science*, volume 765, pages 240–247, 1993.
- [6] John A. Clark and Jeremy L. Jacob. Search for a solution: Engineering tradeoffs and the evolution of provably secure protocols. In *Proceedings of 2000 IEEE Symposium on Research in Security and Privacy*, pages 82–95. IEEE Computer Society, May 2000.
- [7] John A. Clark and Jeremy L. Jacob. Protocols are programs too: the meta-heuristic search for security protocols. *Information and Software Technology*, 43(14):891–904, December 2001.
- [8] Li Gong, Roger Needham, and Raphael Yahalom. Reasoning about belief in cryptographic protocols. In *Proceedings of 1990 IEEE Symposium on Research in Security and Privacy*, pages 234–248. IEEE Computer Society, May 1990.
- [9] D.E. Goldberg. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley, 1989.
- [10] S. Kirkpatrick, C. Gelatt, and M. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, May 1983.
- [11] Adrian Perrig and Dawn Song. A first step towards the automatic generation of security protocols. In *Proceedings of Network and Distributed System Security 2000*, pages 73–83, February 2000.
- [12] Adrian Perrig and Dawn Song. Looking for diamonds in the desert – extending automatic protocol generation to three-party authentication and key agreement protocols. In *Proceedings of The 13th IEEE Computer Security Foundations Workshop*. IEEE Computer Society, July 2000.
- [13] Colin Reeves. *Modern Heuristic Techniques for Combinatorial Problems*. McGraw-Hill Book Company Europe, 1995.
- [14] Paul Syverson and Iliano Cervesato. The logic of authentication protocols. *Foundations of Security Analysis and Design : Tutorial Lectures, Lecture Notes in Computer Science*, 2171:63–136, 2001.
- [15] Paul F. Syverson and Paul C. van Oorschot. A unified cryptographic protocol logic, 1996. NRL Publication 5540–227, Naval Research Lab.