

Cas Cremers



Scyther

User Manual

Draft December 15, 2012

Contents

1	Introduction	5
2	Background	7
3	Installation	9
3.1	Linux	9
3.2	Windows	9
3.3	MAC OS X	9
4	Quick start tutorial	11
5	Input Language	15
5.1	Terms	15
5.1.1	Atomic terms	15
5.1.2	Pairing	15
5.1.3	Symmetric keys	15
5.1.4	Asymmetric keys	16
5.1.5	Hash functions	16
5.1.6	Predefined types	16
5.1.7	Usertypes	16
5.2	Events	17
5.2.1	Receive and send events	17
5.2.2	Claim events and Security properties	17
5.2.3	Internal computation/pattern match events	18
5.3	Role definitions	19
5.4	Protocol definitions	19
5.5	Global declarations	20
5.6	Miscellaneous	20
5.6.1	Macro	20
5.6.2	Include	21
5.6.3	one-role-per-agent	21
5.7	Language BNF	22
5.7.1	Input file	22
5.7.2	Protocols	22
5.7.3	Roles	22
5.7.4	Events	23
5.7.5	Declarations	23
5.7.6	Terms	23

6	Modeling security protocols	25
6.1	Introduction	25
6.2	Example: Needham-Schroeder Public Key	25
7	Specifying security properties	29
7.1	Specifying secrecy	29
7.2	Specifying authentication properties	29
7.2.1	Aliveness	29
7.2.2	Non-injective synchronisation	29
7.2.3	Non-injective agreement	29
7.2.4	Agreement over data	29
8	Using the Scyther tool GUI	31
8.1	Results	31
8.2	Bounding the statespace	33
8.3	Attack graphs	33
8.3.1	Runs	33
8.3.2	Communication events	35
8.3.3	Claims	36
9	Using the Scyther command-line tools	37
10	Advanced topics	39
10.1	Modeling more than one asymmetric key pair	39
10.2	Approximating equational theories	39
10.3	Modeling time-stamps and global counters	41
10.3.1	Modeling global counters	41
10.3.2	Modeling time-stamps using nonces	42
10.3.3	Modeling time-stamps using variables	42
11	Further reading	45
A	Full specification for Needham-Schroeder public key	49

Chapter 1

Introduction

Note: This is a draft of the new version of the Scyther manual. Work on this version started in November 2012, and the manual may therefore be incomplete at points.

Any feedback is welcome and can be sent to Cas Cremers by e-mail: cas.cremers@inf.ethz.ch.

This is the user manual for the Scyther security protocol verification tool.

The purpose of this manual is to explain the details of the Scyther input language, explain how to model basic protocols, and how to effectively use the Scyther tool. This manual does not detail the protocol execution model nor the adversary model used by the tool. It is therefore highly recommended to read the accompanying book that describes, amongst other things, Scyther's underlying protocol model [1].

We proceed in the following way. Some background is given in Chapter 2. Chapter 3 explains how to install the Scyther tool on various platforms. In Chapter 4 we give a brief tutorial using simple examples to show the basics of the tool. Then we discuss things in more detail as we introduce the input language of the tool in Chapter 5. Modeling of basic protocols is described in Chapter 6, and Chapter 7 describes how to specify security properties. The usage of the GUI version of tool is then explained in more detail in Section 8. The underlying command-line tool is described in Section 9. Advanced topics are discussed in Section 10.

Online information

More help can be found online on the Scyther website:

<http://people.inf.ethz.ch/cremersc/scyther/index.html>

Users are advised to subscribe to the Scyther mailing list, whose details can also be found on the Scyther website.

Chapter 2

Background

Scyther is a tool for the formal analysis of security protocols under the *perfect cryptography assumption*, in which it is assumed that all cryptographic functions are perfect: the adversary learns nothing from an encrypted message unless he knows the decryption key. The tool can be used to find problems that arise from the way the protocol is constructed. This problem is undecidable in general, but in practice many protocols can be either proven correct or attacks can be found.

It is not our intention to describe the full protocol model, nor any possible security properties here. For such matters we refer the reader to [1]. Thus, in this manual we assume the reader is familiar with the formal modeling of security protocols and their properties.

Not only is knowledge of security protocol models needed to use the Scyther tool, further knowledge is needed to know how to interpret the results that the tool produces in any useful way. In fact, the reader should be very cautious: security protocol models and their properties are intricate and it is easy to misinterpret the results.

Having said that, one of the main goals of Scyther is to help with the analysis of a protocol in such a way that for example attacks can be understood well. Thus, wherever possible the tool will give useful information on the results.

Chapter 3

Installation

Scyther can be downloaded from the following website:

<http://people.inf.ethz.ch/cremersc/scyther/index.html>

Installation instructions are included. Scyther is available for the Windows, Linux and Mac OS platforms.

3.1 Linux

3.2 Windows

3.3 MAC OS X

Chapter 4

Quick start tutorial

Scyther takes as input a security protocol description that includes security claims, and evaluates these.

Start Scyther by executing the `scyther-gui.py` program in the Scyther directory. The program will launch two windows: the main window, in which files are edited, and the `about` window, which shows some information about the tool.

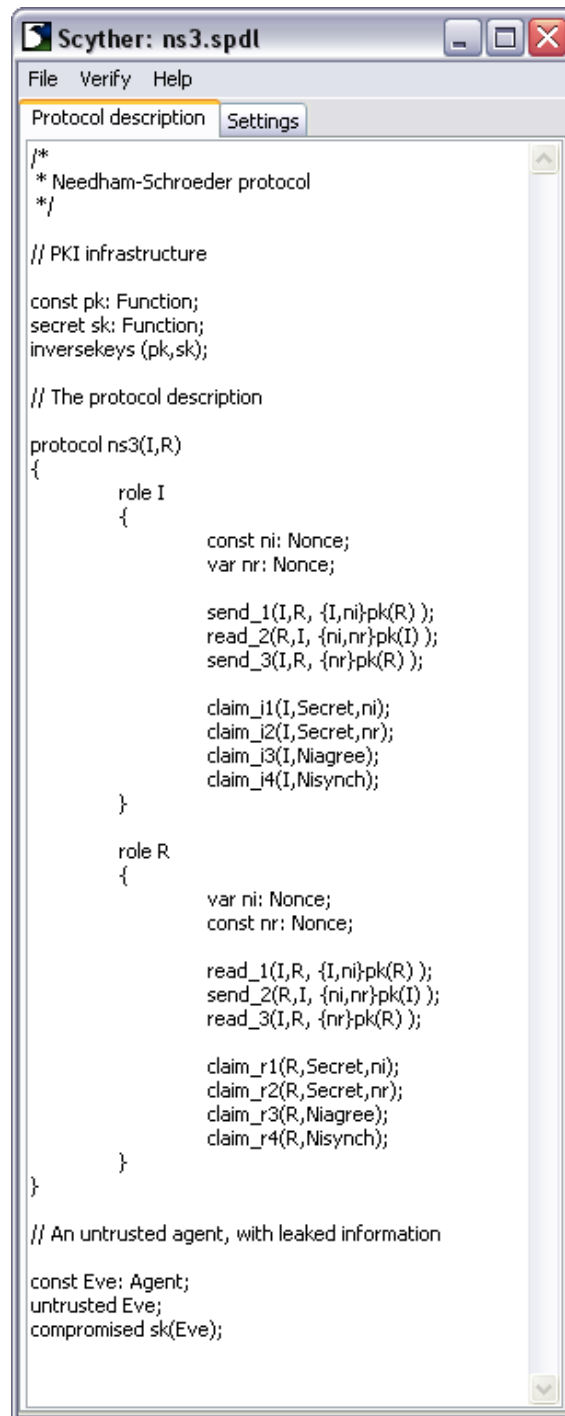
As an introductory example, we will verify the Needham-Schroeder protocol, and investigate an attack on it.

Go to the file→open dialog, and open the file `ns3.spdl` in the Scyther directory. Your main window should look like the one in Figure 4.

By convention, protocol description files have the extension `.spdl` (Security Protocol Description Language), but it can have any name. The file used in this example is shown in Appendix A.

Run the verification tool by selecting `verify→verify_claims` in the menu. A new window will appear during the verification process. Once verification is completed, the window will be replaced by the result window, as shown in Figure 4.

The result window shows a summary of the claims in the protocol, and the verification results. Here one can find whether the protocol is correct, or false. In the next section there will be a full explanation of the possible outcomes of the verification process. Most importantly, if a protocol claim is incorrect, then there exists at least one attack on the protocol. A button is shown next to the claim: press this button to view the attacks on the claim, as in Figure 4.



```

Scyther: ns3.spdl
File Verify Help
Protocol description Settings

/*
 * Needham-Schroeder protocol
 */

// PKI infrastructure

const pk: Function;
secret sk: Function;
inversekeys (pk,sk);

// The protocol description

protocol ns3(I,R)
{
  role I
  {
    const ni: Nonce;
    var nr: Nonce;

    send_1(I,R, {I,ni}pk(R) );
    read_2(R,I, {ni,nr}pk(I) );
    send_3(I,R, {nr}pk(R) );

    claim_1(I,Secret,ni);
    claim_2(I,Secret,nr);
    claim_3(I,Niagree);
    claim_4(I,Nisynch);

  }

  role R
  {
    var ni: Nonce;
    const nr: Nonce;

    read_1(I,R, {I,ni}pk(R) );
    send_2(R,I, {ni,nr}pk(I) );
    read_3(I,R, {nr}pk(R) );

    claim_r1(R,Secret,ni);
    claim_r2(R,Secret,nr);
    claim_r3(R,Niagree);
    claim_r4(R,Nisynch);

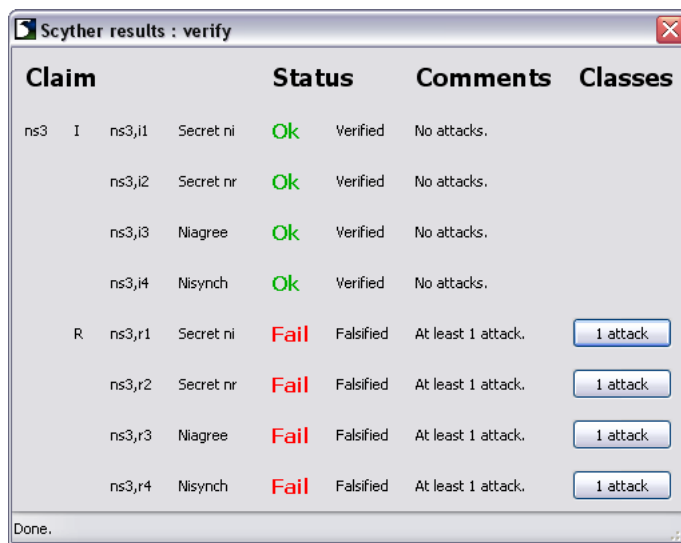
  }
}

// An untrusted agent, with leaked information

const Eve: Agent;
untrusted Eve;
compromised sk(Eve);

```

Figure 4.1: Scyther main window with the file ns3.spdl opened



The image shows a window titled "Scyther results : verify". It contains a table with four columns: Claim, Status, Comments, and Classes. The table lists several claims, some of which are marked as "Fail" with a "1 attack" button next to them. The status "Ok" is in green, and "Fail" is in red. The comments for failed claims are "At least 1 attack.".

Claim	Status	Comments	Classes
ns3 I ns3,i1 Secret ni	Ok	Verified	No attacks.
ns3,i2 Secret nr	Ok	Verified	No attacks.
ns3,i3 Niagree	Ok	Verified	No attacks.
ns3,i4 Nisynch	Ok	Verified	No attacks.
R ns3,r1 Secret ni	Fail	Falsified	At least 1 attack. <input type="button" value="1 attack"/>
ns3,r2 Secret nr	Fail	Falsified	At least 1 attack. <input type="button" value="1 attack"/>
ns3,r3 Niagree	Fail	Falsified	At least 1 attack. <input type="button" value="1 attack"/>
ns3,r4 Nisynch	Fail	Falsified	At least 1 attack. <input type="button" value="1 attack"/>

Done.

Figure 4.2: Scyther result window




Figure 4.3: Scyther attack window

Chapter 5

Input Language

Scyther's input language is loosely based on a C/Java-like syntax. The main purpose of the language is to describe protocols, which are defined by a set of roles. Roles, in turn, are defined by a sequence of events, most of which send or receive terms. We describe these elements in the following sections.

Comments can start with `//` or `#` (for single-line comments) or be enclosed by `/*` and `*/` (for multi-line comments). Note that multi-line comments cannot be nested.

Any whitespace between elements is ignored. It is therefore possible to use whitespace (spaces, tabs, newlines) to improve readability.

A basic identifier consists of a string of characters from the set of alphanumeric characters as well as the symbols `^` and `-`.

The language is case-sensitive, thus `NS3` is not the same identifier as `ns3`.

5.1 Terms

At the most basic level, Scyther manipulates terms.

5.1.1 Atomic terms

An atomic term can be any identifier, which is usually a string of alphanumeric characters.

Atomic terms can be combined into more complex terms by operators such as pairing and encryption.

5.1.2 Pairing

Any two terms can be combined into a term pair: we write (x, y) for the pair of terms `x` and `y`. It is also allowed to write n-tuples as (x, y, z) , which is interpreted by Scyther as $((x, y), z)$.

5.1.3 Symmetric keys

Any term can act as a key for symmetrical encryption.

The encryption of `ni` with a term `kir` is written as:

```
{ ni }kir
```

Unless `kir` is explicitly defined as being part of an asymmetric key pair (explained below), this is interpreted as symmetric encryption.

A symmetric-key infrastructure is predefined: `k(X,Y)` denotes the long-term symmetric key shared between X and Y .

5.1.4 Asymmetric keys

A public-key infrastructure (PKI) is predefined: `sk(X)` denotes the long-term private key of X , and `pk(X)` denotes the corresponding public key.

As an example, consider the following term. It represents the encryption of some term `ni` by the term `pk(I)`. Under normal conventions, this means that the nonce of the initiator (`ni`) is encrypted with the public key of the initiator.

```
{ ni }pk(I)
```

This term can only be decrypted by an agent who knows the secret key `sk(I)`. Section 10.1 describes how to model more than one key pair per agent.

5.1.5 Hash functions

Hash functions are essentially encryptions with a function, of which the inverse is not known by anybody.

They can be used by a global declaration of an identifier to be a hashfunction, e.g.:

```
hashfunction H1;
```

As all agents and protocols should have access to such a function, the declaration of `hashfunction` is usually global, i. e., defined outside of any protocol definition.

Once declared, they can be used in protocol messages, e.g.:

```
H1(ni)
```

5.1.6 Predefined types

Agent Type used for agents.

Function A special type that defines a function term that can take a list of terms as parameter.

By default, it behaves like a hash function: given the term `h(x)` where `h` is of type `Function`, it is impossible to derive `x`.

Nonce A standard type that is often used and therefore defined inside the tool.

Ticket A variable of type `Ticket` can be substituted by any term.

5.1.7 Usertypes

It is possible to define a new type. This can be done using the `usertype` command:


```

usertype MyAtomicMessage;

protocol X(I,R) {
  role I {
    var y: MyAtomicMessage;

    recv_1(I,R, y );
  }
}

```

The effect of such a declaration is that variables of the new type can only be instantiated with messages `m` of that type, i. e., that have been declared by the global declaration `const m: MyAtomicMessage` or the freshly generated `fresh m: MyAtomicMessage` within a role.

In general, the tool can perform better if more is known about which messages might unify or not. By defining a `usertype`, the modeler can inform the tool that a variable can only be instantiated with terms of that type, and not with, e. g., terms of type `Nonce`. Conceptually, one can always write `Ticket` (which corresponds to all possible messages) for each variable type, but then one may find false attacks (in case the implementation in fact does check the type of a message) and the tool will be less likely to verify the property (for an unbounded number of runs).

Draft note (CC) : *Generic (local) declarations? Where is 'fresh' explained?*

Draft note (CC) : *TODO: reverse order: first protocol context, then roles, then events. That way we can actually show what events can do.*

5.2 Events

5.2.1 Receive and send events

The `recv` and `send` events mark receiving and sending a message, respectively.

Note that in some protocol description files one may find the `read` keyword: this is obsolete syntax and can safely be substituted by `recv`.

In most cases, each `send` event will have a corresponding `recv` event. We specify this correspondence by giving such events the same label, denoted by a subscript.

For some protocols we may want to model sending or receiving to the adversary directly, in which case we have no corresponding event. If a `send` or `recv` event has no corresponding event, Scyther will output a warning. To suppress this warning, the label can be prefixed by a bang `!`, e.g.:

```

send_!1(I,I, LeakToAdversary );

```

5.2.2 Claim events and Security properties

Claim events are used in role specifications to model intended security properties. For example, the following claim event models that `sessKey` is meant to be secret.

```

claim(I, Secret, sessKey);

```

There are several predefined claim types.

Secret This claim requires a parameter term. Secrecy of this term is claimed as defined in [1].

Alive Aliveness (of all roles) as defined in [2].

Weakagree Weak agreement (of all roles) as defined in [2].

Commit, Running Non-injective agreement with a role on a set data items [2] can be defined by inserting the appropriate signal claims. In this context, **Commit** marks the effective claim, whose correctness requires the existence of a corresponding **Running** signal in the trace.

These claims are used to model agreement over data, which is explained in Section 7.2.4.

Nisynch Non-injective synchronisation as defined in [1].

Niagree Non-injective agreement as defined in [1].

Reachable When this claim is verified, Scyther will check whether this claim can be reached at all. It is true iff there exists a trace in which this claim occurs. This can be useful to check if there is no obvious error in the protocol specification, and is in fact inserted when the `--check` mode of Scyther is used.

Empty This claim will not be verified, but simply ignored. It is only useful when Scyther is used as a back-end for other verification means. For more on this, see Section 10.

5.2.3 Internal computation/pattern match events

We extend the basic set of events from [1] with two events that can be used to model internal computations.

Match event

New in version v1.1 and Compromise-0.8

The first new event is the `match` event, that is used to specify pattern matching, i. e.,

```
match(pt,m)
```

In operational terms, if there exists a well-typed substitution σ such that $\sigma pt = m$, then this event can be executed. Upon execution, the substitution is applied to the remaining events of the role.

This event can be used to model various constructions, such as equality tests, delayed decryption, checking commitments. They can also be used to model internal computations to simplify specifications, e. g.:

```
var X: Nonce;
var Y;

recv(R,I, X);
match(Y, hash(X,I,R) );
send(I,R, Y,{ Y }sk(I) );
```

In the above example, we could have replaced `Y` by `hash(X,I,R)` throughout the specification, but this version avoid replication.

Not match event**New in version v1.1 and Compromise-0.8**

The second new event is the `not match` event, that is used to specify pattern matching, i. e.,

```
not match(pt,m)
```

The operational interpretation is the opposite of the previous event. If there is no substitution σ such that $\sigma pt = m$, then the event can be executed.

This event can be used to model, e. g., inequality constraints. For example, the execution model allows by default agents executing sessions with themselves. In some cases, we want to exclude such behaviour, because the protocol disallows it. For example,

```
role A {
  not match(A,B);
  send (A,B, m1);
}
```

models a role whose instances only send messages to other agents.

As a more advanced usage, `match` and `not match` can be used together in two roles with a common starting sequence of events to model *if ... then ... else* constructions.

5.3 Role definitions

Role definitions are sequences of events, i. e., declarations, send, receive, or claim events.

```
role Server {
  var x,y,z: Nonce;
  fresh n,m: Nonce;

  send_1(Server,Init, m,n );
  recv_2(Init,Server, x,y, { z }pk(Server) );
}
```

5.4 Protocol definitions

A protocol definition takes as a parameter a sequence of roles, which are then defined within its body.

```
protocol MyProt(Init,Resp,Server)
{
  role Init {
    ...
  }
  role Resp {
    ...
  }
  role Server {
    ...
  }
}
```

```

}
}

```

Helper protocols

It is possible to prepend an “@” symbol before a protocol name. This has no effect on the protocol model, nor on the outcome of the analysis. The “@” is only used when rendering output graphs: the intent is to mark the protocol as a “helper protocol”. Such protocols are often used to model additional adversary capabilities, see Section 10 for examples. When rendering output graphs, Scyther collapses role instances of helper protocols into single nodes. This can make the graphs more readable.

5.5 Global declarations

In many applications global constants are used. These include, for example, string constants, labels, or protocol identifiers.

They are modeled and used in the following way:

```

usertype String;
const HelloWorld: String;

protocol hello(I,R)
{
  role I {
    send_1(I,R, HelloWorld);
  }
  role R {
    recv_1(I,R, HelloWorld);
  }
}

```

5.6 Miscellaneous

5.6.1 Macro

New in version v1.1 and Compromise-0.8

It is possible to define *macros*, i. e., abbreviations for particular term. The syntax used to define these abbreviations is the following:

```

macro MyShortCut = LargeTerm;

```

For example, for a protocol that contains complex messages or repeating elements, macros can be used to simplify the protocol specification:

```

hashfunction h;

protocol macro-example-one(I,R) {
  role I {
    fresh nI: Nonce;

```

```

    macro m1 = h(I,ni);

    send_1(I,R, { m1 }pk(R) );
    claim(I, Secret, m1);
}
role R {
    var X: Ticket;

    recv_1(I,R, { X }pk(R) );
}
}

```

Note that macros have *global scope*, and are handled at the *syntactical* level. This also allows for global abbreviations of protocol messages, e.g.:

```

hashfunction h;
macro m1 = { I,R, nI, h(nI,R) }pk(R);

protocol macro-example-two(I,R) {
    role I {
        fresh nI: Nonce;

        send_1(I,R, m1 );
    }
    role R {
        var nI: Nonce;

        recv_1(I,R, m1 );
    }
}
}

```

Note that in the above example, *nI* is a freshly generated nonce in the I role, and a variable in the R role. Because the macro definitions are unfolded syntactically, the same macro can be used to refer to both terms.

5.6.2 Include

It is possible to import other files in a protocol specification:

```
include "filename";
```

where *filename* denotes the name of the file that will be included at this point. Using this command, it is possible to share e.g. a set of common definitions between files. Typically this will include definitions for the key structures, and (untrusted) agent names. Nested use of this command is possible.

5.6.3 one-role-per-agent

New in version v1.1 and Compromise-0.8

The operational semantics allow agents to perform any roles, and even multiple different roles in parallel. This modeling choice corresponds to the worst possible scenario, in which the adversary

has the most options to exploit. However, in many concrete settings, agents perform only one role. For example, the set of servers may be disjoint from the set of clients, or the set of RFID tags may be disjoint from the set of readers. In such cases, we do not need to consider attacks that exploit that an agent can perform multiple roles. This can be modeled by the following statement:

```
option "--one-role-per-agent"; // disallow agents in multiple roles
```

This causes Scyther to ignore attacks in which agents perform multiple roles. Phrased differently, this corresponds to the situation in which each role is performed by a dedicated set of agents.

5.7 Language BNF

The full BNF grammar for the input language is given below. In the strict language definition, there are no claim terms such as `Niagree` and `Nisynch`, and neither are there any predefined type classes such as `Agent`. Instead, they are predefined constant terms in the Scyther tool itself.

5.7.1 Input file

An input file is simply a list of `spdl` constructions, which are global declarations or protocol descriptions.

$$\langle spdlcomplete \rangle ::= \langle spdl \rangle \{ \text{'\;'} \langle spdl \rangle \}$$

$$\langle spdl \rangle ::= \langle globaldeclaration \rangle \\ | \langle protocol \rangle$$

5.7.2 Protocols

A protocol is simply a container for a set of roles. Because we use a role-based approach to describing roles, the protocol container in fact only affects the naming of the roles: a role “I” in a protocol “ns3” will internally be assigned the name “ns3.I”. This is used to make role names globally unique.

$$\langle protocol \rangle ::= \text{'protocol'} \langle id \rangle \text{'('} \langle termlist \rangle \text{'\;')} \text{'{'} \langle roles \rangle \text{'\;'} [\text{'\;'}]$$

5.7.3 Roles

$$\langle roles \rangle ::= \langle role \rangle [\langle roles \rangle] \\ | \langle declaration \rangle [\langle roles \rangle]$$

$$\langle role \rangle ::= [\text{'singular'}] \text{'role'} \langle id \rangle \text{'{'} \langle roledef \rangle \text{'\;'} [\text{'\;'}]$$

$$\langle roledef \rangle ::= \langle event \rangle [\langle roledef \rangle] \\ | \langle declaration \rangle [\langle roledef \rangle]$$

5.7.4 Events

```

<event> ::= 'recv' <label> '(' <from> ',' <to> ',' <termlist> ')' ';'
| 'send' <label> '(' <from> ',' <to> ',' <termlist> ')' ';'
| 'claim' [ <label> ] '(' <from> ',' <claim> [ ',' <termlist> ] ')' ';'

```

```

<label> ::= '_' <term>

```

```

<from> ::= <id>

```

```

<to> ::= <id>

```

```

<claim> ::= <id>

```

5.7.5 Declarations

```

<globaldeclaration> ::= <declaration>
| 'untrusted' <termlist> ';'
| 'usertype' <termlist> ';'

```

```

<declaration> ::= [ 'secret' ] 'const' <termlist> [ ':' <type> ] ';'
| [ 'secret' ] 'fresh' <termlist> [ ':' <typelist> ] ';'
| [ 'secret' ] 'var' <termlist> [ ':' <typelist> ] ';'
| 'secret' <termlist> [ <type> ] ';'
| 'inversekeys' '(' <term> ',' <term> ')' ';'
| 'compromised' <termlist> ';'

```

```

<type> ::= <id>

```

```

<typelist> ::= <type> { ',' <type> }

```

5.7.6 Terms

```

<term> ::= <id>
| '{' <termlist> '}' <key>
| '(' <termlist> ')'
| <id> '(' <termlist> ')'

```

```

<key> ::= <term>

```

```

<termlist> ::= <term> { ',' <term> }

```


Chapter 6

Modeling security protocols

6.1 Introduction

The correct modeling of a security protocol for analysis in the Scyther tool requires a basic understanding of the underlying symbolic model. This model is explained in detail in [1].

Roughly speaking, the symbolic analysis focuses on the following aspects:

- Logical message components and their intended function within the protocol (public versus secret, freshly generated in each run or constant)
- Message structure (pairing, encryption, signing, hashing)
- Message flow (order, involved agents)

Many other elements are abstracted away. For example, bit strings are abstracted into terms, bit strings that occur with negligible probability are abstracted away, and more complex control flow constructs such as loops are often unfolded for a (low) finite number of times.

6.2 Example: Needham-Schroeder Public Key

As an example, we show how to model a simple protocol.

Figure 6.1 depicts the Needham-Schroeder Public Key protocol. For simplicity, we have only displayed the claim by each role that the initiator nonce ni is secret.

We start off the protocol description by adding a multi-line comment that describes the protocol and other interesting details. Multi-line comments start with `/*` and end with `*/`.

```
1  /*  
2   * Needham-Schroeder protocol  
3   */
```

The protocol uses the default public/private key infrastructure: an agent A has a key pair $(pk(A), sk(A))$.

The protocol has two roles: the initiator role I and the responder role R . We also add a single line comment, starting with `//`.

```

5 // The protocol description
6
7 protocol ns3(I,R)
8 {

```

Scyther works with a role-based description of the protocols. We first model the initiator role. This role has two values that are local to the role: the nonce that is created by I and the nonce that is received. We have to declare them both.

```

9   role I
10  {
11    fresh ni: Nonce;
12    var nr: Nonce;

```

We now model the communication behaviour of the protocol. Needham-Schroeder has three messages, and the initiator role sends the first and last of these. Note the labels (e.g. `_1`) at the end of the `send` and `recv` keywords: these serve merely to retain the information of the connected arrows in the message sequence chart.

```

14   send_1(I,R, {I,ni}pk(R) );
15   recv_2(R,I, {ni,nr}pk(I) );
16   send_3(I,R, {nr}pk(R) );

```

Draft note (CC) : *This should be refactored, and moved to the next chapter.*

Finally, we add the security requirements of the protocol. Without such claims, Scyther does not know¹ what needs to be checked.

Here we have chosen to check for secrecy of the generated and received nonce, and will check for non-injective agreement and synchronisation.

```

18   claim_i1(I,Secret,ni);
19   claim_i2(I,Secret,nr);
20   claim_i3(I,Niagree);
21   claim_i4(I,Nisynch);
22 }

```

This completes the specification of the initiator role.

For this simple protocol, the responder role is very similar to the initiator role². In fact, there are only a few differences:

1. The keywords `var` and `fresh` have swapped places: `ni` was created by I and a freshly generated value there, but for the role R it is the received value and thus a variable.
2. The keywords `send` and `recv` have swapped places.
3. The claims should have unique labels, so they have changed, and the role executing the claim is now R instead of I.


¹If you are unsure about the claims, you can also use the `--auto-claims` switch to automatically generate these at run-time.

²In general, the transformation is not that simple, but for many protocols this will suffice.

The complete role description for the responder looks like this:

```
24  role R
25  {
26    var ni: Nonce;
27    fresh nr: Nonce;
28
29    recv_1(I,R, {I,ni}pk(R) );
30    send_2(R,I, {ni,nr}pk(I) );
31    recv_3(I,R, {nr}pk(R) );
32
33    claim_r1(R,Secret,ni);
34    claim_r2(R,Secret,nr);
35    claim_r3(R,Niagree);
36    claim_r4(R,Nisynch);
37  }
38 }
```

The full protocol description file for the *Needham-Schroeder* protocol is given in Appendix A.



Chapter 7

Specifying security properties

7.1 Specifying secrecy

7.2 Specifying authentication properties

7.2.1 Aliveness

7.2.2 Non-injective synchronisation

7.2.3 Non-injective agreement

7.2.4 Agreement over data

In order to specify data agreement, e. g., that the role I agrees with the role R on a set of terms, e. g., the nonces ni and nr , one inserts two claims:

1. At the end of the I role, insert `claim(I,Commit,R,ni,nr)`;
2. In the R , just before the last send (in case of a protocol with multiple roles: the last send that causally precedes the claim in the I role), insert `claim(R,Running,I,ni,nr)`;

For an example of the use of these claims, see the “ns3.spdl” input file in the Scyther distribution. For a formal definition of the signals, see [2].

Chapter 8

Using the Scyther tool GUI

The Scyther tool can be used in two main ways. First, through the graphical user interface (GUI) and second, through the command-line interface. For most users the first option is preferred.

In this section we detail the Scyther output when used through the GUI.

8.1 Results

As shown before, verifying the Needham-Schroeder public key protocol yields the following results as in Figure 8.1.

The interpretation is as follows: all the claims of the initiator role `ns3,I` are correct for an unbounded number of runs.

Unfortunately, all the claims of the responder role are false. Scyther reports that it found at least one attack for each of those four claims. We could choose to view these attacks: this will be shown in Section 8.3.

In the result window, Scyther will output a single line for each claim. The line is divided into several columns. The first column shows the protocol in which the claim occurs, and the second shows the role. In the third column a unique claim identifier is shown, of the form `p,l`, where `p` is the protocol and `l` is the claim label.¹ The fourth column displays the claim type and the claim parameter.

Under the header `Status` we find two columns. The fifth column gives the actual result of the verification process: it will yield `Fail` when the claim is false, and `Ok` when the claim is correct. The sixth column refines the previous statement: in some cases, the Scyther verification process is not complete (which will be explored in more detail in the next section). If this column states `Verified`, then the claim is provably true. If the column states `Falsified`, then the claim is provably false. If the column is empty, then the statement of fail/ok depends on the specific bounds setting.

The seventh column, `Comments`, serves to explain the status of the results further. In particular, the column contains a single sentences. We describe the possible results below.

- `At least X attack(s)`

Some attacks were found in the state space: however, due to the undecidability of the problem, or because of the branch and bound structure of the search, we cannot be sure that there are no other attack states.

¹This includes the protocol name, which is important when doing multi-protocol analysis.

Claim				Status	Comments	Classes
ns3	I	ns3,i1	Secret ni	Ok	Verified	No attacks.
		ns3,i2	Secret nr	Ok	Verified	No attacks.
		ns3,i3	Niagree	Ok	Verified	No attacks.
		ns3,i4	Nisynch	Ok	Verified	No attacks.
	R	ns3,r1	Secret ni	Fail	Falsified	At least 1 attack. <input type="button" value="1 attack"/>
		ns3,r2	Secret nr	Fail	Falsified	At least 1 attack. <input type="button" value="1 attack"/>
		ns3,r3	Niagree	Fail	Falsified	At least 1 attack. <input type="button" value="1 attack"/>
		ns3,r4	Nisynch	Fail	Falsified	At least 1 attack. <input type="button" value="1 attack"/>

Done.

Figure 8.1: Scyther results for the Needham-Schroeder protocol

In the default setup, Scyther will stop the verification process after an attack is found.

- **Exactly X attack(s)**

Within the statespace, there are exactly this many attacks, and no others.

- **At least X pattern(s)**

- **Exactly X pattern(s)**

These correspond exactly to the previous two, but occur in case of a ‘Reachable’ claim. Thus, the states that are found are not really attacks but classes of reachable states.

- **No attacks within bounds**

No attack was found within the bounded statespace, but there can possibly be an attack outside the bounded statespace.

- **No attacks**

No attack was found within the (bounded or unbounded) statespace, and a proof can be constructed that there is no attack even when the statespace is unbounded. Thus, the security property has been successfully verified.

Note that because of the nature of the algorithm, this result can even be obtained when the statespace is bounded.

8.2 Bounding the statespace

During the verification process, the Scyther tool explores a proof tree that covers all possible protocol behaviours. The default setting is to *bound* the size of this tree in some way, ensuring that the verification procedure terminates. However, importantly, even if the size of this proof tree is bounded, unbounded verification may still be achieved.

In most cases, the verification procedure will terminate and return results before ever reaching the bound. However, if the verification procedure reaches the bound, this is reported in the result window, e.g.:

No attack within bounds

This should be interpreted as: Scyther did not find any attacks, but because it reached the bound, it did not explore the full tree, and it is possible that there are still attacks on the protocol.

The default way of bounding the *maximum number of runs*, or protocol instances. This can be changed in the **Settings** tab of the main window. If the maximum number of runs is e.g. 5, and Scyther reports **No attack within bounds**, this means that there exist no attacks that involve 5 runs or less. However, there might exist attacks that involve 6 runs or more.

For some protocols, increasing the maximum number of runs can lead to complete results (i.e. finding an attack or being sure that there is no attack), but for other protocols the result will always be **No attack within bounds**.

Note that the verification time usually grows exponentially with respect to the maximum number of runs.

8.3 Attack graphs

In Figure 8.3 we show an attack window in more detail.

The basic elements are arrows and several kinds of boxes. The arrows in the graph represent ordering constraints (caused by the prefix-closedness of events in the protocol roles, or by dependencies in the intruder knowledge). The boxes represent creation of a run, communication events of a run, and claim events.

8.3.1 Runs

Each vertical axis represents a run (an instance of a protocol role). Thus, in this attack we see that there are two runs involved. Each run starts with a diamond shaped box. This represents the creation of a run, and is used to give information about the run.

For the run on the left-hand side in the attack we have this information:

<pre>Run #1 Agent2 in role I I -> Agent2 R -> Agent1</pre>
--

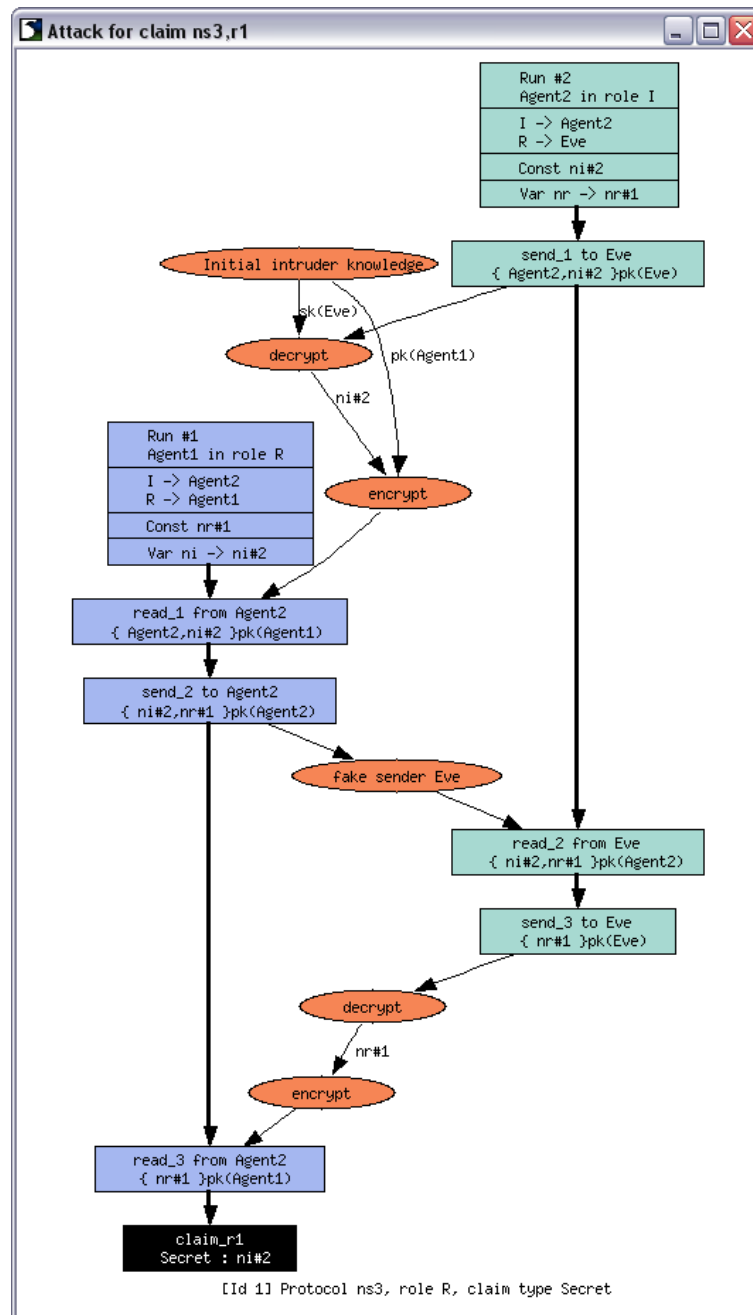


Figure 8.2: Scyther attack window

Each run is assigned a run identifier (here 1), which is an arbitrary number that enables us to uniquely identify each run. This run executes the R role of the protocol. It is being executed by an agent called *Agent1*, who thinks he is talking to *Agent2*. Note that although run 2 is being

executed by **Agent2**, this agent does not believe he is talking to **Agent1**.

```
Run #2
Agent2 in role I
I -> Agent2
R -> Eve
```

In the run on the right, we see This run represents an instance of the role **I**. From the second line we can see which agent is executing the run, and who he thinks he is talking to. In this example, the run is executed by an agent called **Agent2**, who thinks the responder role is being executed by the untrusted agent **Eve**.²

Additionally, the run headers contain information on the freshly generated values (e.g. run 1 generates **nr#1**) and information on the instantiation of the local variables (e.g. run 1 instantiates its variable **ni** with the nonce **ni#2** or run 2).

8.3.2 Communication events

Send events denote the sending of a message. The first send occurs in this attack is the first send event of run 2.

```
send_1(Eve, { Agent#0, ni#2 }pk(Eve) )
```

Every time a message is sent, it is effectively given to the intruder. In this case, because the intruder knows the secret key **sk(Eve)** of the agent **Eve**, he can decrypt the message and learns the value of the nonce **ni#2**.

Receive events correspond to the successful reception of a message. The first receive event that can occur in this attack is the first **recv** event of run 0.

```
recv_1(Agent#0, { Agent#0, ni#2 }pk(Agent#1) )
```

This tells us that the agent executing this run, **Agent#1**, reads a message that is apparently coming from **Agent#1**. The message that is received is **{ Agent#0, ni#2 }pk(Agent#1)**: the name of the agent he thinks he is communicating with and the nonce **ni#2**, encrypted with his public key.

The incoming arrow does not indicate a direct sending of the message. Rather, it denotes an ordering constraint: this message can only be received *after* something else has happened. In this case, we see that the message can only be received after run 2 sends his initial message. The reason for this is the nonce **ni#2**: the intruder cannot predict this nonce, and thus has to wait until run 2 has generated it.

In the graph the connecting arrow is red and has a label “construct” with it: this is caused by the fact that the message sent does not correspond to the message that is received. We know the intruder can only construct the message to be received after the sent message, and thus it must be the case that he uses information from the sent message to construct the message that is received. Other possibilities include a green and a yellow arrow. A yellow arrow indicates that a message was sent, and received in exactly the same form: however, the agents disagree about who was sending a message to whom. It is therefore labeled with “redirect” because the intruder

²Because this agent is talking to the untrusted agent, of course all information is leaked, and no guarantees can be given.

must have redirected the message. A green arrow (not in the picture) indicating that a message is received exactly the same as it was sent, representing a normal message communication between two agents.

Note that a `recv` event without an incoming arrow denotes that a term is received that can be generated from the initial knowledge of the intruder. There is no such event in the example, but this can occur often. For example, if a role reads a plain message containing only an agent name, the intruder can generate the term from his initial knowledge.

8.3.3 Claims

Chapter 9

Using the Scyther command-line tools

All of the features offered by the Scyther GUI are also available through command-line tools. Additionally, the command-line tools offer some features that currently cannot be accessed through the GUI.

Chapter 10

Advanced topics

10.1 Modeling more than one asymmetric key pair

Asymmetric keys are typically modeled as two functions: one function that maps the agents to their public keys, and another function that maps agents to their secret keys.

By default, each agent x has a public/private key pair $(\mathbf{pk}(x), \mathbf{sk}(x))$.

To model other asymmetric keys, we first define the two functions, which are for example named `pk2` for the public key function, and `sk2` for the secret key function.

```
const pk2: Function;  
secret sk2: Function;
```

We also declare that these functions represent asymmetric key pairs:

```
inversekeys (pk2,sk2);
```

If defined in this way, a term encrypted with `pk2(x)` can only be decrypted with `sk2(x)` and vice versa.

10.2 Approximating equational theories

The operational semantics underlying Scyther currently only consider syntactic equality: two (ground) terms are equal if and only if they are syntactically equivalent. However, there are several common cryptographic constructions that are more naturally modeled by using certain equalities. For example:

1. $g^{ab} \pmod N$ and $g^{ba} \pmod N$, to model Diffie-Hellman exponentiation.
2. $k(A, B)$ and $k(B, A)$, to model bidirectional long-term keys.

Although Scyther does not provide direct support for such equational theories, there exists a straightforward underapproximation.

The core idea is that instead of modeling the term equality, we provide the adversary with the ability to learn all terms in an equivalence class if he learns one of its elements. For example, for the equivalence class $\{k(A, B), k(B, A)\}$ we can provide the adversary with the ability to learn $k(B, A)$ from $k(A, B)$, and vice versa. We can model this by introducing an appropriate helper protocol (denoted by the prefix '@'):

```

protocol @keysymmNaive(X) {
  role X {
    var Y: Agent;

    recv_!1(X,X, k(X,Y) );
    send_!2(X,X, k(Y,X) );
  }
}

```

Because the role can be instantiated for any agents X and Y , this covers all possible combinations of agents.

The above naive approximation can be significantly improved. One obvious and practically relevant omission is that the adversary usually learns encrypted messages, but not the key. In such cases, we still would like to model that $\{ m \}_k(A,B) = \{ m \}_k(B,A)$. Thus we adapt our helper protocol:

```

protocol @keysymmInefficient(X,Y) {
  role X {
    var Y: Agent;

    recv_!1(X,X, k(X,Y) );
    send_!2(X,X, k(Y,X) );
  }
  role Y {
    var X: Agent;
    var m: Ticket;

    recv_!1(Y,Y, { m }_k(X,Y) );
    send_!2(Y,Y, { m }_k(Y,X) );
  }
}

```

If the protocol contains further terms in which the symmetric keys appear in other positions, such as in nested encryptions or hashes, we would add further roles.

Draft note (CC) : *Maybe add pointer to more elaborate example in appendix?*

The above approximation is often inefficient in practice. We can improve performance by making the helper protocol rules more tight, i. e., by exploiting more type information about the protocol. For example, if the protocol transmits two types of encrypted messages:

1. $\{ I, nI, nR \}_k(I,R)$, and
2. $\{ nI \}_k(I,R)$,

then we would modify the helper protocol in the following way:

```

protocol @keysymm(X,Y,Z) {
  role X {
    var Y: Agent;

```



```

    recv_!1(X,X, k(X,Y) );
    send_!2(X,X, k(Y,X) );
  }
  role Y {
    var X,Z: Agent;
    var n1,n2: Nonce;

    recv_!1(Y,Y, { Z,n1,n2 }k(X,Y) );
    send_!2(Y,Y, { Z,n1,n2 }k(Y,X) );
  }
  role Z {
    var X,Y: Agent;
    var n1: Nonce;

    recv_!1(Z,Z, { n1 }k(X,Y) );
    send_!2(Z,Z, { n1 }k(Y,X) );
  }
}

```

In general, one would manually inspect the protocol and extract all positions in which a term from an equivalence class occurs as a subterm. For each of these positions, we model an appropriate role in the helper protocols.

This is also used to model, for example, Diffie-Hellman exponentiation. For exponentiation we introduce an abstract function symbol, e. g., `exp`, and a public constant `g`. We then introduce a helper protocol with roles to model that $\text{exp}(\text{exp}(g,X),Y) = \text{exp}(\text{exp}(g,Y),X)$.

In practice, this type of underapproximation has proven to be extremely effective, to the point that all known attacks on real-world protocols that can be modeled using the “real” equational theory, are found by Scyther when using the underapproximation.

One caveat is that while this approximation works well for secrecy and data-agreement, it can cause message-based agreement properties (such as synchronisation) to fail, because their message equality checks are syntactical. These checks are not affected by the introduction of helper protocols.

10.3 Modeling time-stamps and global counters

Scyther’s underlying protocol model currently does not provide support for variables that are shared among the runs of an agent. Effectively, each run starts with a “clean slate”, independent of any runs that have been executed previously. In other words, globally update state can not be modeled directly.

In the following sections we provide some modeling approaches for common problems.

10.3.1 Modeling global counters

Globally incremented counters can be modeled using freshly generated values. This ensures that each run uses a different value. The model is coarse in the sense that the recipient of such a counter cannot check that it is the successor of the previous value of the counter.

10.3.2 Modeling time-stamps using nonces

There are at least two ways to model time-stamps.

The first model is more appropriate for protocols where the probability that a given time-stamp value is accepted by two runs is very low. This occurs when time-stamps have great precision or when two runs occur only sequentially, possibly with some delay time in between. In this case, one can model time-stamps as freshly generated values, e. g., nonces. To cater for the fact that the adversary typically knows the time (and thus can also predict time-stamps), we prepend a send event to the role that provides the adversary with the value of the time-stamp that will be used. For example, we would prepend the send with label !T1 for time-stamp T1 as in the following example:

```

usertype Timestamp;

protocol MyProtocol(Server,Client) {
  role Server{
    fresh T1: Timestamp;

    /* Time-stamps are unique per run */
    send_!T1(Server, Server, T1);

    ...
    /* Server uses time-stamp value */
    send_2(Server,Client, { Server, T1 }pk(Client) );
    ...
  }
}

```

10.3.3 Modeling time-stamps using variables

The second model is more appropriate when it is reasonable that two runs may accept the same time-stamp value. This is common for coarse time-stamps, or for roles that are typically executed with high parallelism, such as server roles. In such cases, one can instead model timestamps as values that are determined by the adversary. In contrast to the previous solution, this is done by prepending a receive event. For example:

```

usertype Timestamp;

protocol MyProtocol(Server,Client) {
  role Server{
    var T1: Timestamp;

    /* Adversary chooses time-stamp value */
    recv_!T1(Server, Server, T1);

    ...
    /* Server uses time-stamp value */
    send_2(Server,Client, { Server, T1 }pk(Client) );
    ...
  }
}

```

```
}  
}
```


Chapter 11

Further reading

Bibliography

- [1] Cas Cremers and Sjouke Mauw. *Operational Semantics and Verification of Security Protocols*. Information Security and Cryptography. Springer, 2012.
- [2] G. Lowe. A hierarchy of authentication specifications. In *Proc. 10th IEEE Computer Security Foundations Workshop (CSFW)*, pages 31–44. IEEE, 1997.

Appendix A

Full specification for Needham-Schroeder public key

```
/*
 * Needham-Schroeder protocol
 */

// The protocol description

protocol ns3(I,R)
{
  role I
  {
    fresh ni: Nonce;
    var nr: Nonce;

    send_1(I,R, {I,ni}pk(R) );
    recv_2(R,I, {ni,nr}pk(I) );
    claim(I,Running,R,ni,nr);
    send_3(I,R, {nr}pk(R) );

    claim_i1(I,Secret,ni);
    claim_i2(I,Secret,nr);
    claim_i3(I,Alive);
    claim_i4(I,Weakagree);
    claim_i5(I,Commit,R,ni,nr);
    claim_i6(I,Niagree);
    claim_i7(I,Nisynch);
  }

  role R
  {
    var ni: Nonce;
```

```
    fresh nr: Nonce;

    recv_1(I,R, {I,ni}pk(R) );
    claim(R,Running,I,ni,nr);
    send_2(R,I, {ni,nr}pk(I) );
    recv_3(I,R, {nr}pk(R) );

    claim_r1(R,Secret,ni);
    claim_r2(R,Secret,nr);
    claim_r3(R,Alive);
    claim_r4(R,Weakagree);
    claim_r5(R,Commit,I,ni,nr);
    claim_r6(R,Niagree);
    claim_r7(R,Nisynch);
  }
}
```

Index

- @, 20
- !, 17
- abbreviate, 20
- Agent**, 16
- agreement (on messages), 18
- agreement on data, 18, 29
- Alive**, 18
- asymmetric key
 - multiple pairs per agent, 39
- asymmetric keys, 16
- at least X attack(s), 31
- at least X pattern(s), 32
- atomic term, 15
- attack graph, 33
- attack window, 33
- bidirectional keys, 39
- BNF, 22
- case-sensitive, 15
- claim**, 17
- claim event, 17
- command-line tools, 37
- comments, 15
- Commit**, 18, 29
- communication event, 35
- const**, 20
- construct, 35
- counter, *see* global counter
- data agreement, 18, 29
- define macro, 20
- Diffie-Hellman exponentiation, 39
- downloading Scyther, 9
- Empty**, 18
- equational theories, 39
- event
 - claim, 17
 - match, 18
 - not match, 19
 - recv, 17
 - send, 17
- events, 17
- exactly X attack(s), 32
- exactly X pattern(s), 32
- exponentiation, *see* Diffie-Hellman exponentiation
- Function**, 16
 - global counter, 41
 - global declarations, 20
 - GUI, 31
 - using Scyther without GUI, 37
- hash functions, 16
- hashfunction**, 16
- helper protocol, **20**, 39
- identifier, 15
- import file, 21
- include**, 21
- input file, 21
- installing Scyther, 9
- internal computation events, 18
- $k(X, Y)$, 16
- macro**, 20
- match event, 18
- message agreement, 18
- multiple asymmetric key pairs per agent, 39
- multiple roles per agent, 21
- Needham-Schroeder protocol, 25
- Niagree**, 18
- Nisynch**, 18
- no attacks, 32
- no attacks within bounds, 32
- non-injective agreement, 18, 29
- non-injective synchronisation, 18

Nonce, 16
not match event, 19
NS, *see* Needham-Schroeder protocol

one role per agent, 21
one-role-per-agent, 21

pairing, 15
pattern match events, 18
pk(X), 16
protocol definition, 19

quick start tutorial, 11

Reachable, 18
read, 17
recv, 17
role definition, 19
run, 33
Running, 18, 29

Scyther website, 5
Secret, 17
security properties, 17
send, 17
sk(X), 16
symmetric keys, 15
synchronisation, 18
syntactic equality, 39

Ticket, 16
time-stamps, 41–43
tupling, 15

usertype, 16

verification, 32

Weakagree, 18
website, *see* Scyther website
whitespace, 15