

Modelling mobility aspects of security policies

Pieter Hartel, Pascal van Eck, Sandro Etalle, Roel Wieringa *

January 26, 2004

Abstract

Security policies are rules that constrain the behaviour of a system. Different, largely unrelated sets of rules typically govern the physical and logical worlds. However, increased hardware and software mobility forces us to consider those rules in an integrated fashion. We present SPIN models of four case studies where mobility plays a role. In each case the model captures both the system of interest and its security policy. The model is then formally checked against a property that represents a principle from the problem domain. The model checking activity shows many examples of policies that are too weak to cope with mobility.

1 Introduction

Security policies are important both in the physical and the logical world; the problem is that they have hardly been studied in an integrated fashion. By a *policy* we mean “a rule that defines a choice in the behaviour of a system” [2]. A *security* policy rules out behaviour “that has been deemed unacceptable” [8]. A *spatial* security policy constrains this further by ruling out behaviour tied to particular locations [9]. Our view on policies is broader than that of other authors because we consider not only the logical world, but also the physical world.

The urgency of solving the “integration problem” is caused by increased mobility. For example (small) mobile computing equipment storing confidential company data is easily carried out of a building, past unsuspecting security guards, regardless of any measures like encryption, or tamper resistance. Mobility thus ties logical and physical security closely together, causing new and, as we will show, unanticipated security problems to arise. We believe that methods and tools are needed to understand and solve the integration problem. We propose a first step towards such a method, with initial tool support. The method is based on developing a formal model of a system with integrated security policies, and on using a model checker to analyse the models. From the analysis we are able to predict the occurrence of new security

problems that would not have occurred without mobility. We present case studies from four different domains to illustrate the method.

To provide structure on the modelling activity, we cast our models in the form of a *policy pattern*. By a *pattern* we mean a class of problems coupled with a partial class of solutions. The class of (integration) problems that we address can be characterised as follows: by introducing mobility, hitherto appropriate security policies become insufficient. The class of solutions consists of two steps. The first step identifies a security *principle*, which is a general guideline for the design of system security. The second step of the solution translates the principle into an abstract specification for the system and its security policy. Since we are dealing not just with the logical world, our principles come from the logical world as well as the physical world. (The term policy pattern has been used before in the sense of a conventional design pattern for Mobile Java code [6].)

We express a policy pattern directly in Promela, the input language of the SPIN model checker [3]. A policy pattern is general; it essentially separates and relates three relevant aspects, viz. the System of interest, its security Policy and a relevant security Principle.

A System of interest is characterised by a trace of relevant events, for example system calls, or messages across a network. This is modelled in a most general way to capture only the essence of the system, particularly its mobility aspects. The system is modelled in SPIN by global data, channels and processes.

A Security Policy constrains the behaviour of the system to acceptable behaviour. The policy constrains traces of events using the same terms as the system, e.g. system calls, network messages etc. The policy is modelled by one or more processes.

A Security Principle is a design guideline for the system and the policy, for example “be reluctant to trust”. To make the principle operational we translate it into a specification for the system and its policy, which, in our case studies, takes the form of a SPIN trace declaration.

The system, policy and principle are formally related

*Dept. of Computer Science, Univ. of Twente, email: {pieter,vaneck,etalle,roelw}@cs.utwente.nl

as follows:

$$(system \parallel policy) \models principle$$

In practice we use the SPIN model checker to analyse a system and policy with respect to a principle; a system and policy that does not abide by the principle gives rise to concrete counter examples. We demonstrate the approach by presenting four case studies in subsequent sections. In each case we identify system, policy, and principle. SPIN then shows that no errors are found by model checking $(system \parallel policy)$ against built-in properties, such as absence of deadlocks and assertion violations. However, model checking $(system \parallel policy) \models principle$ gives traces showing that the principle can be violated. In some of the case studies the violations can be directly attributed to mobility. In other case studies the violations themselves are not caused by mobility but there the case study would not make sense without mobility.

Our modelling methodology is based on a combination of techniques developed by Alan Mycroft and his PhD students from Cambridge, UK, for dynamic security policies in the logical domain (security policy abstraction [5]) and the physical domain (spatial security policy [9]). We extend the Cambridge work in the logical domain [5] by adding mobility considerations. Both works from Cambridge propose policy languages but no realised tool support; we suggest using SPIN by way of tool support. The inspiration for our policy pattern comes from the work of Cheng et al [1] on security patterns, who also use SPIN to model check security policies, however without consideration of either mobility or the physical domain.

Our contribution is that we show how to analyse a system and its security policy while taking both physical and local aspects into account. Such an analysis shows that mobility gives rise to unexpected violations of the security policy, thus contributing to a better understanding of the system.

The four subsequent sections discuss one case study each. The final section concludes and discusses further work.

2 Ping

The first case study concerns the UNIX utility ping, which sends an IP packet to a network node, reporting on the availability of the destination and on the performance of the connection. Making the connection requires root permission (for the socket call), which is potentially dangerous and should thus be minimised. This is captured by the principle of the least privilege, which states that [7]: “Every program and every user of the system should operate using the least set of privileges necessary to complete the job”. In the case of ping (and other commands) the prin-

ciple translates into *dropping root privilege after executing the first socket system call*.

2.1 The System

We model behaviour by traces of system calls, which is abstract because we ignore the parameters of the system calls and any associated calculations. At the same time, naming system calls is concrete because we distinguish many system calls that do not need root permission, and which could be collapsed. The enumeration type `mtype` below mentions the same system calls as used by Mycroft et al [5]. We assume the presence of two different kinds of environment, one in which tests take place, and where ping is permitted to operate. The second environment is a production environment in which ping is prohibited.

```
mtype = {
    Test_Env, Production_Env,
    LibC_exit, LibC_gethostbyname,
    LibC_gettimeofday, LibC_printf,
    SysCall_break, SysCall_mprotect,
    SysCall_recvfrom, SysCall_sendto,
    SysCall_sigaction, SysCall_socket
} ;
```

We use a synchronous channel to connect the ping system to the ping policy. A message consists of a header, i.e. one of the symbols `LibC_exit ... SysCall_socket` and a body, which may either be `Test_Env` or `Production_Env`.

```
chan c = [0] of {mtype, mtype} ;
```

The mobility between environments is modelled by the global variable `env`, and the process `mobility` below, which non-deterministically chooses between environments.

```
byte env = Test_Env ;

active proctype mobility() {
    do
        :: env = Test_Env
        :: env = Production_Env
    od
}
```

The ping system is modelled abstractly by the process `ping_system`. The behaviour generated encompasses all possible traces of the ten system calls in the enumeration type `mtype`. This represents considerably more behaviour than real ping commands would exhibit and includes for example hacked versions of ping. This gives us a good model of reality. Each message is adorned by the environment in which ping is currently performing its duty (through the `env` parameter). Parallel composition of `mobility` with `ping_system` models ping with

extreme mobility, because the value of `env` may change from system call to system call. This is not a realistic model of current systems but it might be realistic for future systems. Imagine a tiny system with the size of a dust particle and a slow processor ¹. The time taken for the system to be blown out of the window may be similar to the time taken to execute a system call, making our extreme mobility example reality!

```
active proctype ping_system() {
  do
    :: c!LibC_exit(env)
    :: c!LibC_gethostbyname(env)
    :: c!LibC_gettimeofday(env)
    :: c!LibC_printf(env)
    :: c!SysCall_break(env)
    :: c!SysCall_mprotect(env)
    :: c!SysCall_recvfrom(env)
    :: c!SysCall_sendto(env)
    :: c!SysCall_sigaction(env)
    :: c!SysCall_socket(env)
  od
}
```

2.2 The Policy

The ping policy below represents an abstract version of the security policy described by Mycroft et al [5]. (We have abstracted away from the fact that extra `break` and `printf` system calls are always allowed.)

We define CPP macros corresponding to the language constructs of the same name proposed by Mycroft et al [5]. (The backward slashes at the end of each line except the last ensure that the macro definition extends across multiple lines.)

```
#define optional( x ) \
  if \
  :: x \
  :: skip \
  fi

#define multiple( x ) \
  do \
  :: x \
  :: break \
  od
```

The first non-deterministic choice below represents a call to ping that responds with a usage message. The second choice represents ping doing its proper work, i.e. a socket call in a test environment, optionally followed by a call to `LibC_gethostbyname` etc. The security policy prohibits socket calls in production environments, all other system calls may occur in any environment.

¹<http://robotics.eecs.berkeley.edu/~pister/SmartDust/>

```
active proctype ping_policy() {
  mtype env ;
  do
    :: c?LibC_printf(env) ;
    c?LibC_exit(env)
    :: c?SysCall_socket(Test_Env) ;
    optional(c?LibC_gethostbyname(env)) ;
    c?LibC_printf(env) ;
    multiple(c?SysCall_sigaction(env)) ;
    multiple(c?SysCall_sendto(env) ; \
             c?SysCall_recvfrom(env) ; \
             optional(c?SysCall_break(env))) ;
  od
}
```

To understand how the ping system and the policy interact, compare the send actions `c!...` of `ping_system` to the receive actions `c?...` of `ping_policy`. This comparison reveals that the former is prepared to engage in any send action, whereas the latter is prepared only to engage in specific receive actions. This then explains how the policy constrains the generic behaviour of the system. We use the same technique throughout the paper to effectuate the security policies.

The separation of system and policy thus provides a convenient way to talk about the system (which is general) and the policy (which constrains general behaviour to specific behaviour). To indicate that this not an entirely trivial result we point out that many other ping implementations are possible that cannot be constrained to conform to the policy, for example a ping implementation where the `do ... od` would be replaced by `if ... fi`.

2.3 The Principle

The ping policy must satisfy the principle of the least privilege. This is mentioned but not explicitly specified as such by Mycroft et al [5]. As stated before, the principle is interpreted as ping must drop root privilege after one socket call, which we interpret here by saying that one socket call is ok, but not two. This is captured by the trace declaration below, which matches a trace that has exactly one system call. When model checking, SPIN tries to find a sample trace that does not match the trace declaration, which is then a counter example for the desired principle.

```
#define anything_but_socket() \
  c?LibC_exit(_) \
  :: c?LibC_gethostbyname(_) \
  :: c?LibC_gettimeofday(_) \
  :: c?LibC_printf(_) \
  :: c?SysCall_break(_) \
  :: c?SysCall_mprotect(_) \
  :: c?SysCall_recvfrom(_) \
  :: c?SysCall_sendto(_) \
  :: c?SysCall_sigaction(_)
```

```

trace {
  do
  :: anything_but_socket()
  :: c?SysCall_socket(Test_Env) -> break
  od ;
  do
  :: anything_but_socket()
  od
}

```

2.4 Analysis

Model checking the parallel composition of `ping_system`, `mobility` and `ping_policy` against the built in formulae of the SPIN model checker (absence of deadlock) shows no errors. This demonstrates that the model is indeed a sensible one.

Model checking the system and the policy against the principle reveals traces that do not match the trace declaration, i.e. traces that violate the principle. A concrete example is `SysCall_socket,Test_Env, LibC_printf,Test_Env, SysCall_socket,Test_Env`. To remedy the situation we have several options. For example we could replace `do ... od` in the ping policy by `if ... fi`, because then only one socket call would result. A better alternative would be to exit the loop once the second non-deterministic choice has completed, because this allows an arbitrary number of ‘usage’ message to be generated but one ‘proper’ ping call.

3 Database application

The second case study investigates the separation of testing and production environments in a modern central database application. The application consists of three layers: data, business logic and presentation. The data and business logic layers reside on a central server. There are two datasets: one with production data and one with randomised test data. The business logic layer accepts network connections from the presentation layer, which consists of Java applets available throughout the organisation. Upon accepting a connection from a presentation layer client, the business logic layer should check the physical location of the client in a configuration database and depending on this location, use either the production or test dataset. We treat the presentation layer and the data base layer as the system, and the business logic layer as the policy because the business logic layer decides what constitutes acceptable behaviour.

3.1 The System

We define the symbols necessary for our example. `Test_0` represents the test data, `Data_0 ... Data_3` represent production data. As in the previous case study, `Test_Env` and `Production_Env` identify the test and production environments. The symbols `Connect ... Reply` represent commands exchanged between the three layers of the system.

```

mtype = {

```

```

  Test_0, Data_1, Data_2, Data_3,
  Test_Env, Production_Env,
  Connect, Disconnect, Request, Reply
}

```

The model comprises three processes (representing the presentation, business logic and database layers) and two synchronous channels connecting the layers. `p2b` connects the presentation to the business logic layer and `b2d` connects the business logic layer to the database layer. The message header may be `Connect ... Reply`. Depending on the message header, the message body may be one of `Test_0 ... Data_3` or `Test_Env` or `Production_Env`.

```

chan p2b = [0] of {mtype, mtype} ;
chan b2d = [0] of {mtype, mtype} ;

```

The presentation layer process below generates an almost arbitrary sequence of requests for both production and test environments. The only form of protocol obeyed is that the presentation layer insists that it receives a reply after each request.

```

active proctype presentation_layer() {
  mtype data ;
end:
  do
  :: p2b!Connect(Test_Env)
  :: p2b!Connect(Production_Env)
  :: p2b!Request(Test_Env) ->
    p2b?Reply(data)
  :: p2b!Request(Production_Env) ->
    p2b?Reply(data)
  :: p2b!Disconnect(Test_Env)
  :: p2b!Disconnect(Production_Env)
  od
}

```

The database layer reports test or production data as appropriate.

```

active proctype database_layer() {
  mtype data ;
end:
  do
  :: b2d?Request(Production_Env) ->
    if
    :: data = Data_1 ;
    :: data = Data_2 ;
    :: data = Data_3
    fi ;
    b2d!Reply(data)
  :: b2d?Request(Test_Env) ->
    b2d!Reply(Test_0)
  od
}

```

3.2 The Policy

The business logic layer process mediates between the presentation and the database layers, ensuring behaviour consistent

with the business rules. In particular the business logic layer insists that applications make data base requests only when connected to the data base. Each request is passed on to the database layer, indicating whether to use the test or production data base.

```
active proctype business_logic_layer() {
  mtype env, data ;
end:
do
  :: p2b?Connect(env) ->
  do
    :: p2b?Request(env) ->
    b2d!Request(env) ;
    b2d?Reply(eval(data)) ;
    p2b!Reply(data)
  :: p2b?Disconnect(env) ->
  break
  od
od
}
```

3.3 The Principle

The principle of interest is that “Testing and production environments are physically separated”. This is enforced by checking that once connected from a particular environment, all following request messages emanate from that same environment until a disconnect message arrives, again from the same environment. An implementation that checks the location once for each connection suffices in the case of a fixed network. In the case of a mobile network that supports roaming, this does not suffice.

```
trace {
  do
  :: p2b?Connect(Test_Env) ->
  do
    :: p2b?Request(Test_Env) ->
    p2b?Reply(_)
  :: p2b?Disconnect(Test_Env) ->
  break
  od
  :: p2b?Connect(Production_Env) ->
  do
    :: p2b?Request(Production_Env) ->
    p2b?Reply(_)
  :: p2b?Disconnect(Production_Env) ->
  break
  od
od
}
```

3.4 Analysis

Model checking the system and the policy against the principle reveals (as expected) that the presentation layer is ill behaved because message sequences with arbitrary test and

production environment parameters are generated. The following trace gives a concrete counter example for our principle: `Connect(Test_Env), Request(Production_Env)`.

The problem lies in the business logic layer (the policy), which should constrain the presentation layer to correct behaviour. Incorporating the principle in the form of a trace declaration points out this deficiency of the business logic layer. The business logic layer might implement the principle using `eval(env)` instead of plain `env` in the receive actions for `Request` and `Disconnect`. This would constrain the environment of the call to match the value of the `env` variable exactly.

4 Smart cards

The third case study investigates the behaviour of next generation smart cards, which are the object of study in the European Inspired project². The principle of interest here is *Be reluctant to trust*. To operationalise this principle we propose security policies that can be customised by the card issuer as well as the card holder. The concrete example that we will study is of (1) a card holder who states that she *does not permit applets to be loaded (i.e. smart card management) other than when she is at the bank* and (2) a card issuer who states that *a payment transaction at a vending machine must always be followed by a loyalty transaction*. Using a smart card that operates under the Be reluctant to trust principle would give the user trust in the system because the user can hold the bank responsible for all the applets on the card. This would also mean that a shop could not load a loyalty applet unless the permission to load applets is explicitly delegated to the shop.

4.1 The System

Following Mycroft et al [5] we assume a tree-shaped world model. Figure 1 shows a shopping mall at the root of the tree, a shop and a bank are at the intermediate layer, and four smart card readers at the leaves.

We use synchronous channels to model the identity of the (physically separated) parties. Three channels connect to the physical locations, four to the smart card readers and another three to the applets. The last three are shown here, the former are defined in the macros `run_leaf` and `run_node` below.

```
chan loyalty = [0] of {chan} ;
chan payment = [0] of {chan} ;
chan management = [0] of {chan} ;
```

We need two kinds of processes: one type representing interior nodes (embedded in macro `run_node`) and one for leaf nodes (macro `run_leaf`). (The macros generate both a channel with the given name, for example `vending`, and corresponding process `proc_vending`.) To model mobility, each process allows an applet to travel from any of its input channels to any of its output channels. In the node process this means that an applet can be received either from the parent of the node, or from one of the children. Similarly, the applet can be moved on to the parent or one of the children.

²<http://www.inspiredproject.com>




Figure 1: Mall with a vending machine and a point of sale terminal (pos) at the shopw, and an automated teller machine (atm) and the traditional counter at the bank.

```

#define run_node(parent, left, right) \
chan parent = [0] of {chan} ; \
active proctype proc_/**/parent() { \
  chan applet ; \
end: \
do \
  :: if \
  :: parent?applet \
  :: left?applet \
  :: right?applet \
  fi ; \
  if \
  :: parent!applet \
  :: left!applet \
  :: right!applet \
  fi \
od \
}

```

A leaf process can exchange applets with the parent only. An applet can in principle be executed on a smart card connected to a smart card reader located at the leaf. This is modelled by sending the identity of the node onto the applet channel thus: `applet!parent`. A leaf makes a non-deterministic choice whether to execute the applet or not.

```

#define run_leaf(parent) \
chan parent = [0] of {chan} ; \
active proctype proc_/**/parent() { \
  chan applet ; \
end: \
do \
  :: parent?applet ; \
  if \
  :: applet!parent \
  :: skip \
  fi ; \
  parent!applet \
od \
}

```

The world is instantiated to the configuration shown in Figure 1 by the seven macro calls below.

```

run_leaf(vending)
run_leaf(pos)
run_node(shop, vending, pos)
run_leaf(atm)
run_leaf(counter)
run_node(bank, atm, counter)
run_node(mall, shop, bank)

```

The system described above is general. It allows free travel of applets, and is prepared to interact with any applet at any of the nodes. This is too liberal, and we need a policy to constrain the resulting behaviour to acceptable behaviour.

4.2 The policy

The `init` process represents the policy that constrains the behaviour of the system. The `init` process begins by injecting a `loyalty` applet, a `payment` applet and a `management` applet into the system (via the parent channel of the root node `mall`). Eventually the mall will return the three applets, whence the system terminates. During its life time, the `init` process is prepared to receive the identity of the host of any of the applets on the corresponding channels `payment`, `loyalty`, and `management`, indicating that the relevant applet is executed while the smart card is connected to the indicated reader.

```

init {
  chan host ;
  mall!loyalty ;
  mall!payment ;
  mall!management ;
end:
do
  :: loyalty?host
  :: payment?host
  :: management?host
  :: mall?eval(loyalty)
  :: mall?eval(payment)
  :: mall?eval(management)
od
}

```

4.3 The Principle

The trace specification below represents the principle *Be reluctant to trust* as operationalised by the two customised policies: one for the smart card issuer and one for the smart card holder. The first non-deterministic choice below represents a payment transaction at the vending machine that must be followed by a loyalty transaction at the vending machine. The second non-deterministic choice represents that the only possibility for a management transaction to take place is at the counter of the bank. The other non-deterministic choices represent the remaining desirable behaviour.

```

trace {
end:
do
  :: payment?eval(vending) ;
  loyalty?eval(vending)
  :: management?eval(counter)
  :: payment?eval(pos)
  :: payment?eval(atm)
  :: payment?eval(counter)
  :: loyalty?eval(vending)
  :: loyalty?eval(pos)
  :: loyalty?eval(atm)
  :: loyalty?eval(counter)
od
}

```

4.4 Analysis

Model checking reveals (again as expected) that the system is ill behaved because the applets roam freely and therefore execute when the principle prohibits this. A concrete counter example shows that after some preliminaries the management applet travels to the point of sale terminal thus: `mall!management, shop!management, pos!management`. There the management applet executes, which is modelled by sending the identity of the host back to the init process: `management!pos`. The violation of the card holder specific part of the principle can be prevented in the policy by replacing `:: management?host` by `:: magagement?eval(counter)`. It is not easy to enforce also the card issuer specific part of the principle as this links two events (the payment and the loyalty transaction) that could in principle be separated by an arbitrary number of unrelated events. To introduce such linkage into the policy, a notion of history would have to be included, for example by adding a variable to the model. This shows that the separation of principle and policy brings at least notational convenience that would not be available otherwise.

5 Peer to peer music sharing

The last case study concerns a peer to peer music sharing system [4]. The model of the relevant processes, message flows and computations is given in Figure 2. The boxes denote processes and their internal actions, the arrows denote messages exchanged between the processes. We will discuss each of the processes and messages below. The model is abstract in the sense that:

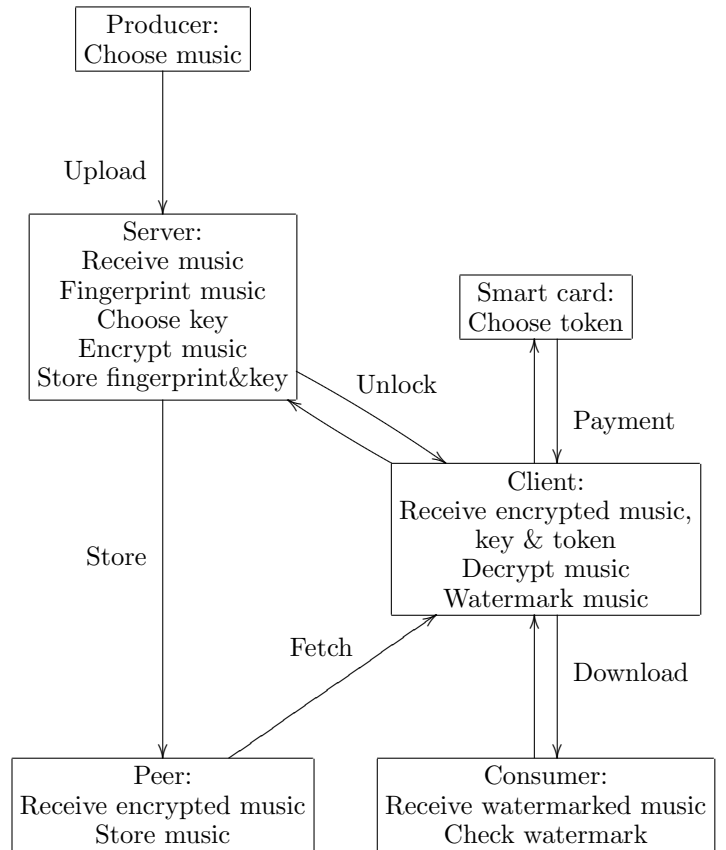


Figure 2: Abstract Music2Share Protocols

- We assume that there are two different classes of users: music (1) producers and (2) consumers; there are valid and expired (3) leases; there are valid and invalid (4) payment tokens; there is (5) music in plain text form, (6) encrypted music, and (7) watermarked music; there are (8) keys and there are (9) fingerprints. All of the above 9 categories are assumed to be distinct and incompatible, for example a fingerprint cannot be confused with the identification of music.
- We assume: (a) a simplistic peer network where peers do not actively redistribute or copy content; (b) the existence of a secure mechanism for looking up the fingerprint for the desired music; (c) idealised encryption, fingerprinting and watermarking.
- We assume that the producer and the server form a secure domain, that the client and the smartcard form another secure domain, and finally that the communication between the client and the server is secure. We make no security assumptions about the peers.
- Without loss of generality, we distinguish precisely two users, two pieces of music, two lengths of lease, two keys etc. This could be extended but no significant new lessons would be learned from doing so.
- We model small numbers of the different parties of the protocol, except the Server, which is centralised. Distributing the server could be accomplished but this would be a refinement that should remain invisible at the chosen level of abstraction.
- We assume synchronous communication so that the network does not have to store messages. We also assume the network to be reliable.

Under the assumptions above we are now able to present the two main scenarios of use: uploading and downloading music.

Scenario 1: upload Starting top left in Figure 2, the producer chooses some music and uploads it onto the server (Upload message). The server receives the music, and calculates the fingerprint that will henceforth identify the music. The server then chooses an encryption key, and encrypts the music. The key is stored with the fingerprint for future use. An appropriate peer stores the encrypted music (Store message) for future reference. At some point in time the server decides that the lease of the music expires and invalidates the key (not shown in the diagram).

Scenario 2a: successful download Starting bottom right in Figure 2, the consumer chooses some music (identified by its fingerprint) and requests the music from a client (Download request). We assume the client receives a valid token from the smartcard by way of payment (Token message). The client then asks the server for the key (Unlock request). We also assume that the lease has not expired so that the client receives a valid key (Unlock reply). The client also receives the encrypted music from the P2P network (Fetch message). The music can now be decrypted and watermarked with the identity of the consumer. The result is sent back to the consumer (Download reply).

Scenario 2b: failed download The scenario will change if either payment could not be arranged (because the valid tokens of the smart card ran out), or when the lease has expired. In both cases the consumer receives an appropriate apology, but no music.

5.1 The System

The relevant symbols of the model are:

```
mtype {
    Download, Fetch, Store,
    Unlock, Payment, Upload,
    Alice, Bob,
    Alice_Music, Bob_Music, No_Music,
    Long_Lease, Short_Lease, Expired_Lease,
    Valid-Token, Invalid-Token,
    Bach, Mozart,
    Bach_Cipher, Mozart_Cipher,
    Bach_Key, Mozart_Key,
    Bach_Fingerprint, Mozart_Fingerprint
}
```

Here Download ... Upload represent the different messages that may be transmitted; Alice, and Bob are the users wishing to download music; Alice_Music, and Bob_Music represent music watermarked by the identity of the user who downloaded the music; No_Music is a place holder for music whose lease has expired; Long_Lease ... Expired_Lease represents three different lengths of lease for shared music; Valid-Token, and Invalid-Token represent two possible payment token values; Bach, and Mozart represent two pieces of music; Bach_Cipher, and Mozart_Cipher represent the encrypted versions of the two pieces of music; Bach_Key, and Mozart_Key represent the encryption keys for the two pieces of music; Bach_Fingerprint, and Mozart_Fingerprint represent the fingerprints of the two pieces of music.

For technical reasons, we need a macro ord (below) to map Mozart_Fingerprint to 0 and Bach_Fingerprint to 1. NIL represents an out-of-band (error) symbol.

```
#define ord(m) (m - Mozart_Fingerprint)
#define NIL    0
```

Keys We assume a unique key for each piece of music. (In addition to CPP macros, SPIN also offers its own variety of macro, the inline):

```
inline choose_key(plain, key) {
    if
    :: plain==Bach ->
        key=Bach_Key
    :: plain==Mozart ->
        key=Mozart_Key
    fi
}
```

Given a plain text, choose_key returns the corresponding key.

Fingerprinting and watermarking Similarly for each piece of music there is a unique fingerprint:

```
inline fingerprint(plain, id) {
  if
  :: plain==Bach ->
    id=Bach_Fingerprint
  :: plain==Mozart ->
    id=Mozart_Fingerprint
  fi
}
```

Once the user has downloaded her Bach or her Mozart, the music is watermarked for her personal use. We are no longer interested in the particular piece of music, only in the user for whom it has been watermarked. An invalid plain text is returned as an invalid marked result `No_Music`.

```
inline watermark(plain, user, marked) {
  if
  :: plain!=NIL && user==Alice ->
    marked=Alice_Music
  :: plain!=NIL && user==Bob ->
    marked=Bob_Music
  :: else ->
    marked=No_Music
  fi
}
```

It should be possible to check whether a piece of content has been watermarked with the correct user identity. `No_Music` does not have a watermark. An incorrect watermark causes a SPIN assertion to fail:

```
inline check_watermark(user, marked) {
  if
  :: user==Alice && marked==Alice_Music ->
    skip
  :: user==Bob && marked==Bob_Music ->
    skip
  :: marked==No_Music ->
    skip
  :: else ->
    assert(false)
  fi
}
```

Encryption and decryption For each piece of music there is one cipher text:

```
inline encrypt(plain, key, cipher) {
  if
  :: plain==Bach && key==Bach_Key ->
    cipher=Bach_Cipher
  :: plain==Mozart && key==Mozart_Key ->
    cipher=Mozart_Cipher
  fi
}
```

With a valid key, the cipher text decrypts uniquely to the original plain text. With an invalid (expired) key, the result is `NIL`:

```
inline decrypt(cipher, key, plain) {
  if
  :: cipher==Bach_Cipher && key==Bach_Key ->
    plain=Bach
  :: cipher==Mozart_Cipher && key==Mozart_Key ->
    plain=Mozart
  :: else ->
    plain=NIL
  fi
}
```

Network The Music2Share network is modelled using two synchronous channels: one for `request` messages and another for `reply` messages. (Channels in SPIN are bi-directional). The channels carry messages with two parameters, where the message header is always one of `Download ... Upload`.

```
chan request=[0] of { mtype, mtype, mtype }
chan reply=[0] of { mtype, mtype, mtype }
```

Server The key server is the only centralised component. The server must create and store keys, it must fingerprint and encrypt music, locate a suitable peer to store encrypted music, and it must serve keys.

Keys are stored together with a lease, which is decremented each time a key is served. This models the process of lease expiry. The data type declaration below defines type `record` holding a key and a lease.

```
typedef record {
  mtype key ;
  mtype lease ;
}
```

The `Server` itself sits in an endless loop waiting for one of two types of messages `Upload` and `Unlock` on the `request` channel.

```
proctype Server() {
  mtype cipher, id, key, plain, user ;
  record store[2] ;
  byte lease ;
  do
  :: request?Upload(plain, lease) ->
    fingerprint(plain, id) ;
    choose_key(plain, key) ;
    store[ord(id)].key = key ;
    store[ord(id)].lease = lease ;
    encrypt(plain, key, cipher) ;
    request!Store(id, cipher)
  :: request?Unlock(id, user) ->
    if
    :: store[ord(id)].lease >
      Expired_Lease ->
        store[ord(id)].lease-- ;
        key = store[ord(id)].key
    :: else ->
      key = NIL
    fi ;
    reply!Unlock(key, user)
  }
```

```

    od
}

```

Upon receipt of an **Upload** message with given **plain** text and **lease**, the server calculates the fingerprint **id**, chooses a **key**, stores the key and the lease in at the appropriate entry **ord(id)** in the array **store**, encrypts the plaintext with the key yielding a **cipher**, and finally transmits a **Store** request onto the network, expecting an appropriate peer to pickup the request and to store the cipher text. This completes the handling of the upload request, no acknowledgement is returned to the requestor of the upload (The network is assumed to be reliable at the chosen level of abstraction).

An **Unlock** request message with a fingerprint **id** and identity **user** causes the server to check the expiry of the lease for the music with the fingerprint **id**. If the lease has expired an invalid key (**NIL**) is created, otherwise the lease is shortened and the correct key retrieved. The key is posted on the **reply** channel, expecting the requestor of the unlock message to pick it up.

Peer A peer is a simple process that serves only to store and communicate the cipher text corresponding to a particular fingerprint.

```

proctype Peer(mtype id) {
  mtype cipher ;
  request?Store(eval(id), cipher) ;
  do
  :: request?Store(eval(id), cipher)
  :: request!Fetch(id, cipher)
  od
}

```

Before the peer enters its main loop, it expects a **Store** message with the initial cipher text. (The expression **eval(id)** states that the actual parameter of the message must have exactly the same value as the variable **id**; the variable **cipher** on the other hand will be bound to what ever actual value is offered by an incoming message). In the main loop, the peer either offers the cipher text to a **Client** in need of the cipher text, or is ready to receive an updated cipher text. If a second process is waiting for a **request!Store** and a third process is waiting for a **request?Fetch**, both transactions are enabled. In this case a non-deterministic choice is made as to which transaction proceeds first.

Smart card The Smart card represents a source of pre-paid tokens.

```

proctype Smartcard() {
  do
  :: request?Payment(_, _) ->
  if
  :: reply!Payment(Valid-Token, NIL)
  :: reply!Payment(Invalid-Token, NIL)
  fi
  od
}

```

The tokens may run out, which is modelled by the **Invalid-Token**. Subsequent valid tokens are the result of recharging the card (not explicitly modelled).

Client The Client mediates between the consumer of music and the Music2Share system.

```

proctype Client() {
  mtype cipher, id, key, plain, marked, user ;
  do
  :: request?Download(id, user) ->
  request!Payment(NIL, NIL) ;
  if
  :: reply?Payment(Valid-Token, _) ;
  request?Fetch(eval(id), cipher) ;
  request!Unlock(id, user) ;
  reply?Unlock(key, eval(user)) ;
  decrypt(cipher, key, plain) ;
  watermark(plain, user, marked) ;
  reply!Download(marked, user)
  :: reply?Payment(Invalid-Token, _) ->
  reply!Download(No_Music, user)
  fi
  od
}

```

The **Client** sits in an endless loop waiting for **Download** messages for a given fingerprint **id** and **user**. The first action is to check payment. If unsuccessful a **NIL** result is returned. Otherwise we **Fetch** the appropriate **cipher** text from a **Peer**. (No request message is necessary here as the peers offer cipher text unsolicited.) Then the client requests the key for the content. The reply message is matched to the identity of the user. After decryption and watermarking the **marked** music is returned to the consumer who posted the download request.

The client receives an invalid key if the lease is expired. In this case the marked result will also be **NIL**.

5.2 The Policy

The producer and the consumer form the endpoints in the value chain, and as such decide the policy for acceptable behaviour. The producer's policy is to upload a choice of music; the consumer downloads a choice of music.

```

proctype Producer() {
  do
  :: request!Upload(Mozart, Long_Lease)
  :: request!Upload(Bach, Short_Lease)
  od
}

```

The **Producer** repeatedly tries to upload **Mozart** (on a long lease) and **Bach**, on a short lease. Further combinations could be added freely.

```

proctype Consumer(mtype user) {
  mtype marked ;
  do
  :: request!Download(Bach_Fingerprint, user) ->

```

```

    reply?Download(marked, eval(user)) ;
    check_watermark(user, marked)
  :: request!Download(Mozart_Fingerprint, user) ->
    reply?Download(marked, eval(user)) ;
    check_watermark(user, marked)
  od
}

```

The **Consumer** does the opposite of the producer: the consumer tries to **Download** content, checking that the downloaded content is indeed for the intended user. The content is identified by its fingerprint; we assume but do not model here the existence of a secure mechanism for looking up the fingerprint for the desired music.

Initialisation The initialisation takes care that all processes are started with the appropriate parameters. Here we choose non-deterministically whether to use Alice or Bob as the consumer.

```

init {
  atomic {
    run Server() ;
    run Peer(Bach_Fingerprint) ;
    run Peer(Mozart_Fingerprint) ;
    run Smartcard() ;
    run Client() ;
    run Producer() ;
    if
      :: run Consumer(Alice) ;
      :: run Consumer(Bob) ;
    fi
  }
}

```

There is one **Server**, for all other processes we assume that at most two versions exist.

5.3 The Principle

The system policy states that she gets the music she has asked for, unless the lease expires, or she fails to pay. This is captured by the `check_watermark` assertion: a failed assertion implies that the policy is violated. This represents acceptable behaviour (from the point of view of the producer) but not desirable behaviour (because the consumer does not get value for money). The guiding principle is thus *value for money*, translated into a trace declaration requiring on the one hand that each time a consumers pays, he or she is guaranteed to get music, and on the other hand that when the customer cannot pay, she gets no music.

```

trace {
  do
    :: reply?Payment(Valid-Token, _) ;
    reply?Unlock(_, _) ;
    if
      :: reply?Download(Alice_Music, _)
      :: reply?Download(Bob_Music, _)
    fi
  }
}

```

```

    :: reply?Payment(Invalid-Token, _) ;
    reply?Download(No_Music, _)
  od
}

```

5.4 Analysis

Model checking the system and the policy reveals that the system does not cause failed assertions, showing that the policy is satisfied. However, the principle may be violated, because the server refuses to deliver an appropriate key once the lease on the music expires. A concrete trace is a little too long to show here; suffice to say that payment takes place, before we request the key. So if the lease expires and no key is forthcoming the user does not get value for money. Swapping the order of payment and key delivery would solve the problem, but at the same time we might introduce a new problem, whereby a key gets delivered for which payment may not be forthcoming. Further study is needed to identify a suitable business policy which would help to decide which alternative is preferred. The point here is that our methodology causes the right questions to be asked during the design stage.

A further point to note is that the trusted computing base (TCB) of the system is small: the producer, client, server and smart card must be secure, but the peers and the consumers do not have to be secure. The peers and the traffic to and from the peers is encrypted by the protocol, and may thus be transported freely on an open network. The music received by the consumer is watermarked with her identity, so that she can play and copy it for her own use, but if she tries to sell it, the watermark will reveal her identity.

6 Conclusions

We are able to model systems, security policies and security principles using a policy pattern that makes these three elements explicit and links them formally thus: $(system \parallel policy) \models principle$.

We have applied the policy pattern to four case studies, showing how often unexpected security problems arise that violate the principle.

In each of the four case studies the system is abstract, the policy is involved but the principle is short and clear. The systems and policies are more difficult to understand because of the concurrency involved. The principles by contrast are not concurrent.

None of the systems satisfy the relevant principle because of the mobility, showing that model checking leads to insight in the case studies.

SPIN's trace declarations are relatively inflexible. It would be useful to either increase the flexibility by changing the SPIN implementation, by generating trace declarations from higher level policies, or a combination of the two approaches.

We are planning to work on a language for policy patterns based on e.g. Ponder [2], from which SPIN models can be generated automatically. This would make it easier for practitioners to use policy patterns. A particular challenge is to relate counter examples generated by the model checker back to the input language.

We also intend to create a library of policy patterns for mobile applications, using the case studies presented here as a starting point.

Finally we should like to investigate ways in which our models of policies and principles can be incorporated in applications by way of execution monitoring.

LNCS 2621, pages 102–117, Warsaw, Poland, Apr 2003. Springer-Verlag, Berlin. <http://www.springerlink.com/link.asp?id=nyxyyrlkbe5c5acc>. 1, 2

References

- [1] B. H. C. Cheng, S. Konrad, L. A. Campbell, and R. Wassermann. Using security patterns to model and analyze security requirements. In C. Hietmeyer and N. Mead, editors, *Int. Workshop on Requirements for High Assurance Systems (RHAS)*, pages 13–22, Monterey, California, Sep 2003. Software Engineering Institute, Carnegie Mellon Univ. <http://www.sei.cmu.edu/community/rhas-workshop/rhas03-proceedings.pdf>. 2
- [2] N. Damianou, N. Dulay, E. Lupu, and M. Sloman. The ponder policy specification language. In M. Sloman, J. Lobo, and E. Lupu, editors, *Int. Workshop on Policies for Distributed Systems and Networks (POLICY)*, volume LNCS 1995, pages 18–38, Bristol, UK, Jan 2001. Springer-Verlag, Berlin. <http://link.springer.de/link/service/series/0558/papers/1995/19950018.pdf>. 1, 11
- [3] G. J. Holzmann. *The SPIN Model Checker: Primer and Reference manual*. Pearson Education Inc, Boston Massachusetts, 2004. 1
- [4] T. Kalker, D. H. J. Epema, P. H. Hartel, R. L. Lagendijk, and M. van Steen. Music2Share - Copyright-Compliant music sharing in P2P systems. *Proceedings of the IEEE Special Issue on Digital Rights Management*, page to appear, 2004. 7
- [5] A. Madhavapeddy, A. Mycroft, D. Scott, and R. Sharp. The case for abstracting security policies. In H. R. Arabnia and Y. Mun, editors, *Int. Conf. on Security and Management (SAM)*, volume 1, pages 156–160, Las Vegas, Nevada, Jun 2003. CSREA Press. <http://cambridgeweb.cambridge.intel-research.net/people/rsharp/publications/sam03-secpol.pdf>. 2, 3, 5
- [6] Q. H. Mahmoud. Security policy: A design pattern for mobile Java code. In *7th. Pattern Languages of Programs Conference (PLoP)*, Allerton Park, Monticello, Illinois, Aug 2000. Washington University. <http://jerry.cs.uiuc.edu/~plop/plop2k/proceedings/Mahmoud/Mahmoud.pdf>. 1
- [7] J. H. Saltzer and M. D. Schroeder. The protection of information in computer systems. *Proceedings of the IEEE*, 63(9):1278–1308, Sep 1975. 2
- [8] F. B. Schneider. Enforceable security policies. *ACM Transactions on Information and System Security*, 3(1):30–50, Feb 2000. <http://doi.acm.org/10.1145/353323.353382>. 1
- [9] D. Scott, A. Beresford, and A. Mycroft. Spatial security policies for mobile agents in a sentient computing environment. In M. Pezzè, editor, *6th Fundamental Approaches to Software Engineering (FASE)*, volume