

Failure Analysis of an E-commerce Protocol using Model Checking*

Indrakshi Ray

Department of Computer and Information Science

University of Michigan-Dearborn

4901 Evergreen Road, Dearborn, MI 48128

Indrajit Ray

Abstract

The rapid growth of electronic commerce (e-commerce) has necessitated the development of e-commerce protocols. These protocols ensure the confidentiality and integrity of information exchanged. In addition, researchers have identified other desirable properties, such as, money atomicity, goods atomicity and validated receipt, that must be satisfied by e-commerce protocols. This paper shows how model checking can be used to obtain an assurance about the existence of these properties in an e-commerce protocol. It is essential that these desirable properties be satisfied, even in the presence of site or communication failure. Using the model checker we evaluate which failures cause the violation of one or more of the properties. The results of the analysis are then used to propose a mechanism that handles the failures to make the protocol failure resilient.

1 Introduction

The growing popularity of e-commerce has resulted in the development of a number of e-commerce protocols. Most of these protocols ensure that the information exchanged between the parties is protected from unauthorized disclosure and modification. Moreover, researchers have identified several other desirable properties of e-commerce protocols. Examples of these properties include money atomicity and goods atomicity [18], and validated receipt [14]. Money atomicity ensures that money is neither created nor destroyed in the course of an e-commerce transaction. Goods atomicity ensures that a merchant receives payment if and only if the customer receives the product. Validated receipt ensures that the customer is able to verify the contents of the product about to be received, before making the payment. Although such properties have been identified, a major problem is verifying if a given e-commerce protocol

satisfies these properties, specially in the presence of network and site failures. In this paper we address the problem of protocol verification using an existing software verification technique known as model checking.

The reasons for using model checking are as follows. First, model checking is a completely automated technique and considerably faster than other approaches, such as, theorem proving [1, 2, 6, 12]. Second, if a property does not hold, a counter example is produced by the model checker which helps in understanding why the property does not hold. Last, but not the least, model checking has previously been used successfully to verify security protocols [8, 9, 10, 11]. In this paper we use the Failure Divergence Refinement (FDR) model checker [7].

An e-commerce protocol, being distributed in nature, is subject to site and/or communications failures. Are the desirable properties satisfied in the event of a failure? Which properties are valid when some site fails? Our formal analysis helps answer these questions. From the analysis we find that, if there is a failure, most properties do not hold for the protocol described in [14]. Thus some extra mechanism needs to be incorporated into the protocol that will ensure the properties even in the face of failure. In this paper we show how to extend the basic protocol to incorporate these mechanisms and make the protocol fault tolerant.

Briefly, our contribution is as follows. First, we show how to model the e-commerce protocol and the desirable properties in a specification language, and then verify these properties using an existing model checker. Formal specification and verification give an increased assurance that the properties are indeed satisfied by the protocol. Second, we show how to model site and communication failures. We use the FDR model checker to identify which failures preserve the properties and which ones do not. Finally, we propose a mechanism to handle the identified failures that do not preserve the desirable properties; integrating this mechanism with the protocol ensures that the resulting protocol satisfies the properties in the presence of failures.

The rest of the paper is organized as follows. Section 2 describes the basic protocol. Section 3 describes some re-

*This work has been partially supported by the National Science Foundation under grant EIA 9977548 and by a Faculty Research Grant from the University of Michigan-Dearborn.

1. $TP \rightarrow C$: *download of encrypted product*
2. $C \rightarrow M$: *purchase order*
3. $M \rightarrow C$: *product encrypted with a second key K_2*
4. $M \rightarrow TP$: *the decrypting key \hat{K} for the product and the approved purchase order*
5. $C \rightarrow TP$: *the payment token and copy of the purchase order*
6. $TP \rightarrow C$: *the decrypting key*
7. $TP \rightarrow M$: *payment token*

Table 1. Messages Exchanged

lated work. Section 4 describes how the basic protocol and the desirable properties can be modeled and verified by the FDR model checker. Section 5 describes how the different forms of failures can be modeled and illustrates which properties get violated by site or communication failures. Section 6 proposes a mechanism that makes the protocol failure resilient. Section 7 concludes the paper.

2 Basic Protocol

We begin with a brief description of the e-commerce protocol proposed by Ray et al. [14] designed for electronic transactions involving digital products. In that work the authors demonstrate informally that their protocol satisfies money atomicity, goods atomicity and validated receipt properties in the absence of failures.

Table 1 shows the different steps in the protocol¹. The protocol works by exchanging messages between a customer (C), a merchant (M) and a trusted third party (TP). The notation $X \rightarrow Y : P$ is used to denote X sends message P to Y . A merchant has several products to sell. The merchant places a description of each product on an on-line catalog service with the third party together with copy of the product encrypted with a key K_1 . If the customer is interested in a product, he downloads the encrypted version of the product (step 1 in table 1) and then sends a purchase order to the merchant (step 2). Note that the customer cannot use the product unless he has decrypted it. Now the merchant does not send the decrypting key unless the merchant receives payment. The customer does not pay unless he is sure that he is getting the right product. This is handled as follows: the merchant sends the product (step 3) encrypted with a second key, K_2 , such that K_2 bears a particular mathematical relation with the key, K_1 . Additionally,

¹Since we do not focus on the security aspects of the protocol, the cryptographic aspects have been abstracted away.

the merchant escrows the decryption key, \hat{K} , corresponding to K_2 , with the third party (step 4). The mathematical relation between the keys K_1 and K_2 , is the basis for the theory of cross validation that has been proposed (see [14] for complete details). Briefly the theory of cross validation states that the encrypted messages compare if and only if the unencrypted messages compare. Thus, by comparing the encrypted product received from the merchant with the encrypted product that the customer downloaded from the third party, the customer can be sure that the product he is about to pay for is indeed the product he wanted. At this stage the customer is yet to obtain the actual product because he does not have the key, \hat{K} , to decrypt the encrypted product. Once the customer is satisfied with his comparison, he sends his payment token to the third party (step 5). The third party verifies the customer's financial information and forwards the decrypting key to the customer (step 6) and the payment token to the merchant (step 7).

3 Related Work

Lowe [9, 10, 11] have used the FDR model checker to find attacks on cryptographic protocols. Roscoe et al. [15] have used the FDR model checker together with data independence techniques to prove that some security protocols are free from attacks.

Heintze et al. [8] focus on the non-security aspects of e-commerce protocols and use the FDR model checker to verify the money and goods atomicity properties of two e-commerce protocols – NetBill [3] and Digicash [4]. The important contribution of the work is that it illustrates how to model the e-commerce protocols and the properties of interest in CSP. This work confirms the claims that the Net-Bill protocol does have money and goods atomicity and provides a counter-example to illustrate why the Digicash protocol does not have money atomicity.

We also use the FDR model checker to analyze an e-commerce protocol. However, we give a more comprehensive treatment of site and communication failures than given by Heintze et al. [8]. We allow the customer, the merchant, the third party, and the communication links to fail arbitrarily. We propose additional mechanisms that ensure that the desirable properties are still preserved in spite of failures.

4 Using FDR to Verify the Properties

FDR [7] is a model checker, based on the theory of CSP, that allows one to check properties of finite state systems. In CSP, processes are described using events and operators. Events cause a process to change state. The representation of states is implicit. A process may be composed of component processes that require synchronization on some events;

each component must be willing to participate in an event before the process can make a state transition. This is how the component processes interact with each other.

In FDR, the finite state system is modeled as a process, say SYSTEM, and the property of interest is modeled as another process, say SPEC. If SYSTEM is a refinement of SPEC (that is, the set of the possible behaviors of SYSTEM is a subset of the set of possible behaviors of SPEC), then the property modeled by SPEC holds. If some property does not hold, FDR generates a counter-example that illustrates under what scenario the property is violated. The counter-example is useful for analyzing and debugging purposes.

4.1 Modeling Communication between Processes

The communication between processes is synchronized in CSP. To model asynchronous communication we follow the approach of Heintze et al. [8]. A sender sends messages over a channel while the receiver receives the message over another channel. Thus when two agents are communicating, two channels are associated with each agent for a total of four channels.

For example, to model the communication between the merchant and the customer four channels, (minc, coutm, cinm, moutc) and two processes (COMMmc and COMMcm) are involved. The process COMMmc reads a data from channel coutm and writes the data to channel minc and the process COMMcm reads a data from channel moutc and writes it to channel cinm. Process COMMcm is modeled in CSP as:

```
COMMcm = []x: {po} @(coutm ?x -> (minc !x -> COMMcm))
```

where {po} – the purchase order – is the set of all data that is communicated directly from the customer to the merchant. Similarly, we have four channels and two processes for modeling the communication between merchant and trusted third party. Please refer to the extended version of this paper [13] for details.

4.2 Modeling the Customer Process

The protocol starts when the customer browses the catalog hosted on the third party and downloads the encrypted product from there. The downloading of the encrypted product is modeled as the sending of the encrypted product by the third party and the receipt of the product by the customer. Thus, we can say that, initially the customer waits for an encrypted product from the third party.

```
CUSTOMER = cint ?x -> DOWNLOADED_EGOODS(x)
```

Once the customer has downloaded the product, it sends a purchase order to the merchant. This is modeled as:

```
DOWNLOADED_EGOODS(x) = coutm !po -> PO_SENT(x)
```

The customer then waits for the encrypted product from the merchant. On receiving a message from the merchant, the customer checks to see if the message is indeed some encrypted product sent by the merchant. If so, the customer proceeds to the next step, otherwise it continues to wait for the encrypted product. The specification for this event is as:

```
PO_SENT(x) = cinm ?y -> if (y==encryptedGoods1 or
    y==encryptedGoods2) then RECEIVED_EGOODS(x,y)
    else PO_SENT(x)
```

The next step involves comparing the encrypted product received from the merchant with those downloaded from the third party. If the two do not match, the customer terminates the protocol.

```
RECEIVED_EGOODS(x,y) = if (x==y) then
    RECEIVED_CORRECT_GOODS else ABORT
```

When the customer is satisfied with the encrypted product, he sends the payment token to the third party.

```
RECEIVED_CORRECT_GOODS =
    coutt !paymentToken -> TOKEN_SENT
```

After sending the payment, the customer waits for a message from the trusted third party. The third party either sends the customer the key or an abort message, depending on the outcome of the protocol. Once the customer has received the message from the third party, the protocol stops. Otherwise the customer continues to wait for the message.

```
TOKEN_SENT = cint ?y ->
    if (y==key) then SUCCESS else
    if (y== transactionAborted) then ABORT
    else TOKEN_SENT
```

4.3 Modeling the Merchant Process

On the merchant side, the protocol begins with the merchant waiting to receive a purchase order from a customer.

```
MERCHANT = minc ?x -> if (x==po) then PO_REC
    else MERCHANT
```

The merchant in response must send an encrypted product to the customer. The merchant can act in two ways – either he sends the correct encrypted product (denoted by encryptedGoods1) or an incorrect encrypted product (denoted by encryptedGoods2). This non-deterministic choice is modeled as follows:

```
PO_REC = (moutc !encryptedGoods1 ->
    ENCRYPTED_GOODS_SENT) |~| (moutc !encryptedGoods2
    -> ENCRYPTED_GOODS_SENT)
```

Once the merchant has sent the encrypted product, he must send the decryption key to the trusted third party.

```
ENCRYPTED_GOODS_SENT = moutt !key -> KEY_SENT
```

After sending the key, the merchant waits to receive the payment token from the third party. The third party either sends the payment token or a transaction abort message if the transaction was aborted. The merchant process terminates once it receives a message, otherwise it continues to wait for the message.

```
KEY_SENT = mint ?x -> if (x==paymentToken)
  then SUCCESS else if (x==transactionAborted)
  then ABORT else KEY_SENT
```

4.4 Modeling the Trusted Third Party Process

The customer downloading the encrypted product, is modeled from the third party's end, as the trusted third party sending the encrypted product to the customer.

```
TP = toutc !encryptedGoods1 -> WAIT_TOKEN_KEY
```

The next step involves the third party waiting to receive the payment token from the customer and the key from the merchant. When the third party receives a message it checks if the message is a payment token or key or neither. Note that, it is not known whether the key or payment token will arrive first. If the payment token arrives first, the third party must wait for the key. On the other hand, if the key arrives first, the third party must wait for the payment token. This aspect of the protocol is modeled as follows:

```
WAIT_TOKEN_KEY = (tinc ?a -> if (a==paymentToken)
  then WAIT_KEY(a) else WAIT_TOKEN_KEY) []
  (tinm ?b -> if (b==key) then WAIT_TOKEN(b)
  else WAIT_TOKEN_KEY)
WAIT_KEY(a) = tinm ?b -> if (b==key) then
  CHECK_TOKEN(a,b) else WAIT_KEY(a)
WAIT_TOKEN(b) = tinc ?a -> if (a==paymentToken) then
  CHECK_TOKEN(a,b) else WAIT_TOKEN(b)
```

Once the third party has received both the key and the payment token, it proceeds to the next step of validating the payment token with the customer's financial institution. The details of the validation process is outside the scope of the protocol and is not modeled. Instead, the model non-deterministically chooses between the options: (i) token okay or (ii) token not okay. If the payment token is okay, the third party proceeds to send out the key to the customer and the token to the merchant. If the payment token is not okay an abort message is sent to the customer and the merchant, and the protocol terminates.

```
CHECK_TOKEN(a,b) = OK_TOKEN(a,b) |~| NOK_TOKEN
OK_TOKEN(a,b) = SEND_TOKEN_KEY(a,b)
NOK_TOKEN = SEND_ABORT_MESSAGE
```

The process of sending an abort message to customer and merchant is modeled by the following step.

```
SEND_ABORT_MESSAGE = toutc !transAborted
  -> toutm !transAborted -> STOP
```

The processes of sending out the key to the customer and the token to merchant can occur in any order. Thus the sending out of key and token by the trusted third party is modeled as follows.

```
SEND_TOKEN_KEY(a,b) = toutc !b -> SEND_TOKEN(a)
  |~| toutm !a -> SEND_KEY(b)
SEND_TOKEN(a) = toutm !a -> SUCCESS
SEND_KEY(b) = toutc !b -> SUCCESS
```

4.5 Modeling the Money Atomicity Property

Our protocol assumes that the merchant receiving the validated token is an acceptable form of payment. The money atomicity property states that money is neither created nor destroyed in the process of an electronic transaction. This property is satisfied when the payment token, sent by the customer, is received by the merchant. However, the payment token may not always be received by the merchant. Consider, for example, the case when the payment token cannot be validated by the third party because of insufficient funds or some other error. In this case, the customer receives a transaction abort message which tells him that the transaction was aborted and the payment token was not forwarded to the merchant. Thus money atomicity is satisfied when one of the following things happen: (i) the customer sends the payment token and the merchant receives it or (ii) the customer sends the payment token and then receives a transaction abort message. This is modeled in CSP as:

```
SPEC1 = STOP |~| ((couth.paymentToken ->
  mint.paymentToken -> STOP) [] (couth.paymentToken
  -> cint.transAborted -> STOP))
```

The next step in the verification of money atomicity involves forming a new process called SYSTEM1 by combining the merchant, the customer and the third party processes with the appropriate communication processes, and hiding all the irrelevant events (that is, all events other than those specified in SPEC1). To verify that SYSTEM1 satisfies SPEC1 we need to check that SYSTEM1 is a failure/divergence refinement [16] of SPEC1. This check is done automatically by FDR, confirming that money atomicity is indeed satisfied by the protocol.

4.6 Modeling the Goods Atomicity Property

The goods atomicity property ensures that a customer receives the product if and only if he has paid. Note that both the correct encrypted product and the key are required by the customer in order to get the product. Thus in our model, the product comprises both the correct encrypted product and the key. The payment is made in the form of a token. The merchant receiving the token suffices as the customer having paid.

In the protocol, the customer may receive the correct encrypted product from the merchant (denoted by `encryptedGoods1`) or some incorrect one (denoted by `encryptedGoods2`). The goods atomicity property requires one of the following things to happen: (i) the customer receives both the correct encrypted product and the keys and the merchant receives the token, or (ii) the customer receives just the encrypted product and neither the merchant gets the payment token nor the customer the keys. For (i), note that the customer may receive the key before the merchant receives the payment token or vice-versa. Thus goods atomicity is specified as:

```
SPEC2 = STOP |~| ((cinm.encryptedGoods1 -> STOP)
[] (cinm.encryptedGoods2 -> STOP) [])
(cinm.encryptedGoods1 -> cint.key ->
mint.paymentToken -> STOP) [] (cinm.encryptedGoods1
-> mint.paymentToken -> cint.key -> STOP))
```

We create another process called `SYSTEM2` by combining the customer, merchant, third party and the communication processes and hiding the unrelated events (the events not associated with the specification of the goods atomicity property). FDR proves that `SPEC2` is a failure divergence refinement of `SYSTEM2` confirming that the protocol does indeed have goods atomicity.

4.7 Modeling the Validated Receipt Property

The validated receipt property ensures that the customer makes the payment only after he is satisfied that the product he is about to receive is the one that he is paying for.

In other words, the validated receipt property ensures one of the following things happen: (i) the customer receives some encrypted product and does not make payment (either because he has received incorrect product or decides not to purchase the product), or (ii) the customer makes the payment after receiving the correct encrypted product. This is modeled as:

```
SPEC3 = STOP |~| ((cinm.encryptedGoods2 -> STOP) []
(cinm.encryptedGoods1 -> STOP) [])
(cinm.encryptedGoods1
-> coutt.paymentToken -> STOP))
```

Next we create a process `SYSTEM3` by combining the processes of the merchant, customer, third party and the communications processes and use FDR to check that `SYSTEM3` is a failure divergence refinement of `SPEC3`. In this manner we are ensured that the protocol also has the validated receipt property.

5 Detecting Violation of Properties due to Failures

An informal analysis reveals that the properties may be violated if the customer, merchant, third party and commu-

nication links fail arbitrarily. The following paragraphs describe how we use the model checker to detect failures that result in the destruction of the properties.

5.1 Introducing Unreliable Communication Channels

The basic protocol assumes that the channels through which the three parties communicate are reliable. That is, any data sent by one party to another will eventually be received by the second party and is not lost in transmission.

In an unreliable channel the data may get lost. That is, data transmitted by the customer may not be received by the merchant. So after the event of the customer transmitting the data, one of two things may happen: (i) the data is lost or (ii) the merchant receives the data. Note that the choice of what is going to happen is not known and happens non-deterministically. The unreliable channel is modeled as:

```
COMMcm = []x: {po} @(coutm ?x ->
(COMMcm |~| (minc !x -> COMMcm)))
```

With this modification, we recheck the properties using FDR. It turns out that this unreliable channel has no effect on the properties. Incrementally we introduce unreliability in the other channels and test for the satisfaction of properties. Our experiments indicate that the properties are violated when the following channels are made unreliable: (i) the channels connecting the third party and the customer and (ii) those connecting the third party to the merchant.

5.2 Introducing Failures in the Customer Process

In the following paragraphs we show how the properties get violated when the customer process fails at certain points.

Consider the first step for the customer process:

```
CUSTOMER = cint ?x -> DOWNLOADED_EGOODS(x)
```

Suppose we allow the customer to abort in this step. The question, then, is does any property get violated? To find out, we need to model the possibility of the customer aborting in the first step:

```
CUSTOMER = ABORT |~| (cint ?x -> DOWNLOADED_EGOODS(x))
```

The above specification says that the customer may abort or wait for the downloading of the encrypted product in a non-deterministic manner. After making the above alteration to the customer process, we use FDR to check for the satisfaction of the properties. As expected, the customer aborting in the first step, has no effect on the properties. We make similar modifications to each step in the customer process and check for the violation of the properties.

Our results indicate that such modification to any step, except in the last step of the customer process (that is after

the customer has sent the payment token), preserves all the properties. Allowing the customer to abort in the last step violates both money atomicity and goods atomicity. To illustrate how the money atomicity property is violated, we use the FDR debugger which generates the following sequence of events.

```
toutc.encryptedGoods1, cint.encryptedGoods1,
coutm.po, minc.po, moutc.encryptedGoods1,
moutt.key, tinm.key, cinm.encryptedGoods1,
coutt.paymentToken, tinc.paymentToken,
toutc.transAborted, toutm.transAborted,
mint.transAborted
```

The above sequence tells us that the following actions are executed. The customer downloads the encrypted product from the third party, then sends a purchase order to the merchant. On receiving the purchase order, the merchant sends the encrypted product to the customer and the key to the third party. The third party receives the key. The customer receives the encrypted product and validates it. The customer then sends the payment token to the third party. At this point, it appears that the customer aborts since we do not see any more messages sent or received by the customer. The third party receives the key from the merchant, the payment token from the customer, and then validates the token. The token turns out to be invalid and an abort message is sent by the third party to the customer and the merchant. Since the customer has aborted in the mean time, he does not get the transaction abort message from the third party. The merchant, however, receives the abort message. In the above scenario, the customer sends out the payment token, but neither the merchant received the payment token nor the customer the transaction abort message. Thus money atomicity is violated. Similarly, a counter example is generated illustrating how goods atomicity was violated. For the sake of brevity, we omit this counter example from the paper.

Thus, our conclusion is that, the customer cannot abort after sending out the payment token and before receiving the key; if the customer does indeed abort we will no longer have money atomicity or goods atomicity.

5.3 Failures in Merchant, Third Party Processes

For lack of space, we do not show similar analysis for the merchant or the third party process. The interested reader is referred to the extended version of this paper [13].

For the merchant process, FDR reported that allowing the merchant to abort in its third step, that is after the merchant has sent the key to the third party results in the violation of money atomicity property. However, a minor modification to the protocol handles this problem. The third party associates a timeout event when waiting for responses from customer or merchant. If the third party is waiting for a

long time, the time out event occurs and it sends an abort message to the customer and merchant.

FDR further reported that allowing the merchant process to abort in the last step, that is, after sending the key but before receiving the payment token, violates both money atomicity and goods atomicity.

Finally, we consider the third party process. The third party process can abort unilaterally only at its first step.

6 Ensuring Failure Resilience of the Protocol

From the above discussion we can summarize: (i) The customer cannot abort after he has sent the payment token to the third party. (ii) The merchant cannot abort after he has sent the product decryption key to the third party. (iii) The third party cannot abort unilaterally after its first step.

To ensure that the e-commerce protocol is resilient to site or link failures we propose the following extension to the basic protocol. We assume that each party involved in the transaction, keeps a copy of the information that it sends to another party – for example purchase order, payment token and so on – in its stable storage till such time as the information is no longer needed. Writes to the stable storage are atomic and durable until intentionally purged.

1. The customer, the merchant and the third party uses a system-wide unique identifier, T_i , to denote the current e-commerce transaction. The identifier is a tuple of the form $\langle PID, C, M \rangle$, where PID is the identifier for the product the customer, C purchases from the merchant, M . The customer stores a log record of the form $\langle T_i, INITIATE \rangle$ to its stable storage and then sends the purchase order to the merchant.
2. When the merchant receives the purchase order, it writes a log record $\langle T_i, INITIATE \rangle$ to its stable storage; then the merchant checks to see if the purchase order is to its satisfaction. If it is not, the merchant writes an abort record in its log – $\langle T_i, ABORT \rangle$ and aborts the transaction. It informs the customer of this decision. Otherwise it sends the encrypted product to the customer and the product decryption key and the approved purchase order to the third party. Finally, it writes a log record to its stable storage of the form $\langle T_i, KEY - SENT \rangle$. At this stage the merchant enters a point of no return; it cannot abort unilaterally.
3. After receiving a message from the merchant the customer checks to see if it is an abort message or the encrypted product. If it is an abort, the customer aborts the transaction and writes a log record of the form $\langle T_i, ABORT \rangle$. Otherwise the customer validates the encrypted product. If validated, the customer sends the payment token and purchase order to the third party

and then writes a log record to its stable storage. The log record is of the form $\langle T_i, PAYMENT - SENT \rangle$. This is the point of no return for the customer. If the encrypted product is not validated the customer can either request the product from the merchant, or abort the transaction.

4. One of the messages - either the message containing the payment token and purchase order from the customer or the message containing the product decryption key and approved purchase order from the merchant - will arrive at the third party before the other message. On receiving the message, the third party associates the unique identifier T_i to this current transaction and writes a log record to its stable storage of the form $\langle T_i, INITIATE \rangle$. The third party starts a timer at this point. If the third party does not receive the other message before the timer expires, it writes a log record $\langle T_i, ABORT \rangle$ and sends abort messages to both the customer and the merchant.
5. After receiving the payment token from the customer, the third party validates the token with the customer's financial institution. If the validation fails the third party writes a log record $\langle T_i, ABORT \rangle$ and informs both the customer and the merchant. Otherwise, after the third party has received both - the product decryption key from the merchant and the payment token from the customer - the third party sends the payment token to the merchant and writes a log record $\langle T_i, PAYMENT - FORWARDED \rangle$, and sends the decryption key to the customer and writes a log record $\langle T_i, KEY - FORWARDED \rangle$.
6. The customer writes the log record $\langle T_i, FINISH \rangle$ after receiving the decryption key from the third party.
7. The merchant also writes a log record $\langle T_i, FINISH \rangle$, after it has received the payment token.

6.1 Protocol Failure Analysis

Let us consider each of the possible failure scenarios individually and see why the protocol is failure resilient. Recall that we are interested in the cases after the customer has sent the payment token or the merchant has sent the product decryption key.

1. **Merchant fails after sending product decryption key but before writing $\langle T_i, KEY - SENT \rangle$.** After recovery from failure the merchant finds from its log that T_i has been initiated but the product decryption key has not been sent out (no information about the key having been sent is recorded). Consequently, it queries the third party to find out the status. If the status is abort,

the merchant aborts. If the third party has not received the key, the merchant resends the key and write the appropriate record. If the third party cannot provide a status, the merchant resends the encrypted product to the customer, and the key and approved purchase order to the third party and writes the appropriate log records. It then waits for the payment token from the third party. Finally, as a result of the status query the merchant may receive the payment token. It then finishes by writing the appropriate log record.

2. **Merchant fails after writing $\langle T_i, KEY - SENT \rangle$ or Merchant fails before writing $\langle T_i, FINISH \rangle$.** After recovery from failure, the merchant finds that it has not received the payment token. It asks the third party for the payment token. The third party responds either by sending the payment token or an abort message. If it is an abort message, the merchant write $\langle T_i, ABORT \rangle$ in its stable storage and aborts. If payment token is received the merchant writes $\langle T_i, FINISH \rangle$ to log.
3. **Customer fails after sending payment token but before writing $\langle T_i, PAYMENT - SENT \rangle$.** After recovery, the customer notes from log that T_i has been initiated but no other information (such as, information about the product received or payment token sent) is recorded in the log. The customer, in this case, gets in touch with the merchant and asks for the product. The merchant either sends the encrypted product or an abort message. If the customer receives the encrypted product, the customer validates it, sends the payment token and writes the appropriate log record.
4. **Customer fails after writing $\langle T_i, PAYMENT - SENT \rangle$ or Customer fails before writing $\langle T_i, FINISH \rangle$.** After recovery the customer notes that the decryption key has not been received. So it requests the third party for the product decryption key. The third part responds with either an abort message or the decryption key. If it is an abort message, the customer writes $\langle T_i, ABORT \rangle$ to its log and aborts. If it is the decryption key, the customer writes $\langle T_i, FINISH \rangle$ to the log and finishes.
5. **Third party fails before writing $\langle T_i, INITIATE \rangle$.** At this stage the third party is not aware of the transaction T_i . Consequently the third party does nothing. At some point of time either the customer or the merchant will get in touch asking for the product decryption key or a status query. At this stage the third party will write the log record $\langle T_i, INITIATE \rangle$ and ask the customer for the payment token and purchase order, and the merchant for the product decryption key and the approved purchase order. It then starts the timer.

6. **Third party fails after writing $\langle T_i, INITIATE \rangle$, or Third party fails before writing either $\langle T_i, PAYMENT - FORWARDED \rangle$ or $\langle T_i, KEY - FORWARDED \rangle$.** After recovery the third party notes that T_i has been initiated. It asks the customer for the payment token and purchase order, and the merchant for the product decryption key and the approved purchase order. Once the third party has received a response, it starts the timer and waits for the other message.
7. **Third party fails after writing one of the records $\langle T_i, PAYMENT - FORWARDED \rangle$ or $\langle T_i, KEY - FORWARDED \rangle$ but before writing the other.** After recovery the third party sends the message that was not sent out and writes the appropriate record.

7 Conclusion and Future Work

In this paper we have illustrated how model checking can be used to get assurance that an e-commerce protocol does satisfy the properties of money atomicity, goods atomicity, and validated receipt. We have also shown how model checking can be used to detect violation of properties in the presence of site and communication failures. Using the analysis, we proposed a mechanism that preserves the properties even in the event of sites or communications failures.

Our experiments indicate that model checking is an extremely powerful technique for protocol verification. If a property does not hold, the model checker generates a counterexample. We use the counterexample to trace the execution of the protocol and determine why and how the property is violated. Often we found that the most obvious specification of a property was incorrectly or inadequately specified. This kind of interactive experimentation helped us expressing the properties and the model more precisely.

A number of issues still remain to be investigated. We plan to investigate the security issues of the e-commerce protocol using model checking. We need to investigate how the many different forms of electronic payment schemes that are available today [5, 17] can be incorporated into our protocol and the resulting solutions be verified for the existence of the properties. We also plan to improve our protocol by reducing the involvement of the trusted third party.

References

- [1] D. Bolognani. An Approach to the Formal Verification of Cryptographic Protocols. In *Proceedings of the 3rd ACM Conference on Computer and Communications Security*, pages 106–118, New Delhi, India, Mar. 1996. ACM Press.
- [2] D. Bolognani. Towards the Formal Verification of Electronic Commerce Protocols. In *Proceedings of the 10th IEEE Computer Security Foundations Workshop*, June 1997.
- [3] B. Cox, J. D. Tygar, and M. Sirbu. NetBill Security and Transaction Protocol. In *Proceedings of the 1st USENIX Workshop in Electronic Commerce*, pages 77–88, July 1995.
- [4] A. F. D. Chaum and M. Naor. Untraceable electronic cash. In *Advances in Cryptology – Proceedings of CRYPTO*, pages 200–212. Springer-Verlag, 1990.
- [5] S. Dukach. SNPP: A Simple Network Payment Protocol. Technical report, MIT Laboratory for Computer Science, 1992.
- [6] B. Dutertre and S. Schneider. Using a PVS Embedding of CSP to Verify Authentication Protocols. In *Theorem Proving in Higher Order Logics*, volume 1275 of *Lecture Notes in Computer Science*, pages 121–136. Springer-Verlag, 1997.
- [7] Formal Systems (Europe) Ltd. *Failure Divergence Refinement - FDR2 User Manual*, version 2.64 edition, Aug. 1999.
- [8] N. Heintze, J. Tygar, J. Wing, and H. Wong. Model Checking Electronic Commerce Protocols. In *Proceedings of the 2nd USENIX Workshop in Electronic Commerce*, pages 146–164, Nov. 1996.
- [9] G. Lowe. Breaking and Fixing the Needham-Schroeder Public-key Protocol Using FDR. In *Tools and Algorithms for the Construction and Analysis of Systems: 2nd International Workshop, TACAS 96*, pages 147–166, 1996.
- [10] G. Lowe. Some New Attacks Upon Security Protocols. In *Proceedings of the IEEE Computer Security Foundations Workshop*. IEEE Computer Society Press, 1996.
- [11] G. Lowe and A. W. Roscoe. Using CSP to detect errors in the TMN protocol. *IEEE Transactions on Software Engineering*, 23:659–669, 1997.
- [12] L. C. Paulson. Proving Properties of Security Protocols by Induction. Technical Report TR409-lcp, Computer Laboratory, University of Cambridge, Dec. 1996.
- [13] I. Ray and I. Ray. Failure Analysis of an E-commerce Protocol using Model Checking. Technical report, University of Michigan-Dearborn, Jan. 2000. available from the URL <http://www.engin.umd.umich.edu/~iray>.
- [14] I. Ray, I. Ray, and N. Narasimhamurthy. A Fair-Exchange Protocol with Automated Dispute Resolution. Technical report, University of Michigan-Dearborn, Jan. 2000. available from the URL <http://www.engin.umd.umich.edu/~iray>.
- [15] A. W. Roscoe. Proving Security Protocols with Model Checkers by Data Independence Techniques. In *Proceedings of the IEEE Computer Security Foundations Workshop*. IEEE Computer Society Press, 1998.
- [16] A. W. Roscoe. *The Theory and Practice of Concurrency*. Prentice Hall, Great Britain, 1998.
- [17] L. H. Stein, E. A. Stefferud, N. S. Borenstein, and M. T. Rose. The Green Commerce Model. Technical report, First Virtual Holdings Incorporated, 1994.
- [18] J. D. Tygar. Atomicity in Electronic Commerce. In *Proceedings of the 15th Annual ACM Symposium on Principles of Distributed Computing*, pages 8–26, May 1996.