

# A Mechanism for Establishing Policies for Electronic Commerce\*

Naftaly H. Minsky<sup>†</sup>

Victoria Ungureanu\*

Department of Computer Science  
Rutgers University

New Brunswick, NJ, 08903 USA

Phone: (908) 445-2085

Fax: (908) 445-0537

Email: {minsky,ungurean}@cs.rutgers.edu

## Abstract

*This paper introduces a mechanism for establishing policies for electronic commerce in a unified and secure manner. A commercial policy can be viewed as the embodiment of a contract between the principals involved in a certain type of commercial activity, and it may be concerned with such issues as: ensuring that a payment for services is refunded under specified circumstances; preventing certificates representing e-cash from being duplicated; ensuring that credit card numbers are used only for the transaction they are intended for; and, for certain socially sensitive transactions like the purchase of drugs, ensuring auditability by proper authorities.*

*Our mechanism is based on a previously published concept of regulated interaction. It makes a strict separation between the formal statement of a policy, which we call a “law,” and the enforcement of this law, which is carried out by a set of policy-independent trusted controllers. A new policy under this scheme is created basically by formulating its law, and can be easily deployed throughout a distributed system. This mechanism enables a single agent to engage in several different activities, subject to disparate policies.*

*Two examples policies are discussed here in detail: one ensures refundability of payment, under certain circumstances; the other provides for payment by means of non-copyable tickets.*

## 1 Introduction

Research on electronic-commerce focused mainly on means for secure and efficient transfer of funds [12, 4,

11]. Such means are indeed necessary for e-commerce, but they are not sufficient. Commercial activities are not limited to simple exchange of funds and merchandise between a client and a vendor—they often consist of multi-step transactions, sometimes involving several participants, that need to be carried out in accordance to a certain policy. A commercial policy is the embodiment of a contract (which may be explicit or implicit) between the principals involved in a certain type of commercial activity—and it may be concerned with such issues as: ensuring that a payment for services is refunded under specified circumstances; preventing certificates representing e-cash from being duplicated; ensuring that credit card numbers are used only for the transaction they are intended for; and, for certain socially sensitive transactions like the purchase of drugs, ensuring auditability by proper authorities.

There is, of course, no universal policy. Many different policies are being employed in traditional commerce, and many will have to be employed in electronic commerce as well. For instance, a recent paper [6] identified several examples of policies for purchasing of information, using such things as *subscription*, *pay-per-view* and *pre-paid voucher*, which use different form of payment, employ different interaction protocols, and are suitable for different circumstances. The potential diversity of e-commerce policies creates a need for convenient and reliable means for defining, deploying, enforcing and using policies. It should, in particular, be possible for a single agent to engage in several different activities that are subject to different policies, as is the case for traditional commerce.

The currently prevailing method for establishing e-commerce policies is to build an interface that implements a desired policy, and distribute this interface

---

\*In The 18th International Conference on Distributed Computing Systems (ICDCS), May 1998.

<sup>†</sup>Work supported in part by NSF grants No. CCR-9626577

among all who may need to operate under it. Unfortunately, such “manual” implementation of e-commerce policies is both unwieldy and unsafe. It is *unwieldy* in that it is time consuming and expensive to carry out, and because the policy being implemented by a given set of interfaces is obscure, being embedded into the code of the interface. A manually implemented policy is *unsafe* because it can be circumvented by any participant in a given commercial transaction, by modifying his interface for the policy. The first difficulty has been addressed recently in [6, 1], and the second has been addressed in [5], but we are not aware of any attempt to alleviate both of these difficulties together.

In this paper we introduce a mechanism, based on what we call “regulated interaction” [9], that can support a wide range of commercial policies in a unified and secure manner. Our mechanism makes a strict separation between the formal statement of a policy, which we call a “law,” and the enforcement of this law, which is carried out by a set of policy-independent trusted *controllers*. A new policy under this scheme is created basically by formulating its law, which is easy to deploy and to use, and quite scalable.

We start, in Section 2, with a definition of our concept of a *policy*, illustrating it with an informal example of a policy that ensures refundability of payments, whose formalization is presented in Section 5. In Section 3 we present the implementation of this concept by a toolkit called Moses, illustrating it with another example policy that provide for payments by means of non-copyable tickets. In Section 4 we explain how the proposed mechanism can be made secure, and we conclude with a thought about some broader implications of this work.

## 2 The Concept of a Policy

Before defining our concept of a policy, we illustrate it by presenting an analogue of the *refund policies* often employed in traditional retail. This analogue is the following “money back” policy, which we denote by  $\mathcal{MB}$ .

Let an order sent to a vendor have the following form `order(specs,fee,gp)`, where `specs` specifies the required merchandise, `fee` is the payment, and `gp` is the “grace period” (say, in minutes) required by the client. It is the essence of this policy that the client be able to get his money back by sending the message `refund(specs)` to the vendor. The refund is granted if either the ordered item has not been shipped, or less than `gp`

minutes have passed since the ordered merchandise arrived at the client<sup>1</sup>.

Note that refund policies are effective in traditional commerce because the clients of large respectable vendors tend to trust them to honor their “money-back guarantee”—since such vendors value their reputation. But this kind of trust does not occur naturally in e-commerce because of the open nature of the Internet. This is why e-commerce policies must be *enforced* to be effective, and why they must be formalized. The following, then, is our definition of the concept of policy for e-commerce.

A policy has three elements: the *actions* being regulated, the *group* of participants that engage in this actions, and the *guiding principles* themselves. More formally:

**Definition 1** A policy  $\mathcal{P}$  is a triple  $(\mathcal{M}, \mathcal{G}, \mathcal{L})$ , where

1.  $\mathcal{M}$  is the set of messages, representing the primitive operations of the activity in question. They are called  $\mathcal{P}$ -messages.
2.  $\mathcal{G}$  is a distributed group of agents, sometimes called a “policy-group,” that are permitted to send and receive  $\mathcal{P}$ -messages, and thus are the participants in policy  $\mathcal{P}$ .
3.  $\mathcal{L}$  is an enforced set of rules—called the law of this policy—regulating the exchange of  $\mathcal{P}$ -messages between members of group  $\mathcal{G}$ .

For example, the components of the  $\mathcal{MB}$  policy are as follows: the group  $\mathcal{G}_{\mathcal{MB}}$  consists of all the vendors and clients participating in it. The set of messages  $\mathcal{M}_{\mathcal{MB}}$  consists of all the messages exchanged during the purchase process, including the messages `order` and `refund` mentioned above. Finally, the law  $\mathcal{L}_{\mathcal{MB}}$  is the set of rules that regulate the exchange of messages, ensuring that clients can get their refund; these rules are to be written in a formal languages to be introduced in the following section.

Several clarifications of our concept of a policy are in order. First, we take a policy to have an independent existence, separate from the agents participating in it, and independent from any other policy. We provide means for an agent to *join* a given policy  $\mathcal{P}$ —subject to the law of this policy—which will enable this agent to send and receive  $\mathcal{P}$ -messages. A given agent may thus join several policies, and be active under them concurrently. Note that by definition, different policies cannot be in conflict because the exchange

<sup>1</sup>The return of the merchandise, if refund is requested, will be discussed later.

of  $\mathcal{P}$ -messages is governed solely by  $\mathcal{L}_{\mathcal{P}}$ , regardless of any other policy the agent in question might belong to.

Second, although a policy can be formulated without detailed knowledge of the behavior of the prospective participants, an agent involved in a certain policy needs to be familiar with its law, and with the structure of its messages. But since the law is enforced, it cannot be violated even by an agent who does not know the law, or who willfully ignores it.

### 3 Regulated Interaction and the Moses Toolkit

We describe here a mode of interaction between distributed agents, under which the exchange of messages between agents is regulated by means of a pre-specified policy, in the sense of Section 2. This mode, called *regulated interaction* (RI), has been implemented by an experimental toolkit called Moses. We start this discussion by introducing the concept of law in detail, illustrating it by an example. In Section 3.3 we introduce our law enforcement mechanism, and in Section 3.4 we describe the manner in which a policy can be deployed throughout the internet.

#### 3.1 The Law of a Policy

A law  $\mathcal{L}$  of a policy  $\mathcal{P}$  determines the treatment of  $\mathcal{P}$ -messages by specifying what should be done when such a message is sent, and when it arrives. More specifically, the law deals with the following two kinds of events that are regulated under Moses<sup>2</sup>:

1.  $\text{sent}(\mathbf{x}, \mathbf{m}, \mathbf{y})$  — occurs when agent  $\mathbf{x}$  sends an  $\mathcal{P}$ -message  $\mathbf{m}$  addressed to  $\mathbf{y}$ . The sender  $\mathbf{x}$  is considered the *home* of this event.
2.  $\text{arrived}(\mathbf{x}, \mathbf{m}, \mathbf{y})$  — occurs when an  $\mathcal{P}$ -message  $\mathbf{m}$  sent by  $\mathbf{x}$  arrives at  $\mathbf{y}$ . The receiver  $\mathbf{y}$  is considered the *home* of this event.

We assume no prior knowledge of, or control over, the occurrence of these *regulated events*. But the effect that any given event  $\mathbf{e}$  would actually have is prescribed by the law  $\mathcal{L}$  of the policy in question. This prescription, called the *ruling* of the law for this event, is a (possibly empty) sequence of *primitive operations* (discussed later) which are to be carried out at the home of  $\mathbf{e}$ , as the immediate response to its occurrence.

Structurally, the law  $\mathcal{L}$  is a pair  $\langle \mathcal{R}, \mathcal{CS} \rangle$ , where  $\mathcal{R}$  is a fixed set of rules defined for the entire group  $\mathcal{G}$  of the policy in question, and  $\mathcal{CS}$  is a mutable set

$\{\mathcal{CS}_x \mid x \in \mathcal{G}\}$  of what we call *control states*, one per member of the group. These two parts of the law are discussed in more detail below.

**The control state  $\mathcal{CS}_x$ .** This is the part of the law  $\mathcal{L}$  that is associated with the individual member  $\mathbf{x}$  of a group. It is a bag of terms, called the *attributes* of this member. The main role of these attributes is to enable  $\mathcal{L}$  to distinguish between different kinds of members, so that the ruling of the law for a given event may depend on its home. Some of the attributes of an agent have a predefined semantics, such as the term  $\text{self}(\mathbf{n})$  that defines  $\mathbf{n}$  to be the unique identifier for this agent (unique within the group). However, the semantics of attributes is defined, in general, by the law, for a given group. For instance, in the implementation of the *money back* policy, described in Section 2, a term  $\text{vendor}$  in the control state of an agent identifies it as one of the vendors.

**The Primitive Operations.** The operations that can be included in the ruling of the law for a given regulated event  $\mathbf{e}$ , to be carried out at the home of this event, are called *primitive operations*. They are “primitive” in the sense that they can be performed *only* if thus authorized by the law. These operations include:

1. Operations that change the CS of the home agent. Specifically, we can perform the following operations: (1)  $+\mathbf{t}$  which adds the term  $\mathbf{t}$  to the control state; (2)  $-\mathbf{t}$  which removes the term  $\mathbf{t}$ ; (3)  $\mathbf{t}_1 \leftarrow \mathbf{t}_2$  which changes term  $\mathbf{t}_1$  with term  $\mathbf{t}_2$ ; and (4)  $\text{incr}(\mathbf{t}(\mathbf{v}), \mathbf{x})$  which increments the value  $\mathbf{v}$  of a term  $\mathbf{t}$  with some quantity  $\mathbf{x}$ .
2. Operation  $\text{forward}(\mathbf{x}, \mathbf{m}, \mathbf{y})$  emits to the network the message  $\mathbf{m}$  addressed to  $\mathbf{y}$ , where  $\mathbf{x}$  identifies the *sender* of the message. (When a message thus forwarded to  $\mathbf{y}$  arrives, it would trigger at  $\mathbf{y}$  the event  $\text{arrived}(\mathbf{x}, \mathbf{m}, \mathbf{y})$ .) The most common use of this operation is in a ruling for event  $\text{sent}(\mathbf{x}, \mathbf{m}, \mathbf{y})$ , where operation  $\text{forward}$  (with no arguments) simply completes the passing of the intended message.
3. Operation  $\text{deliver}(\mathbf{x}, \mathbf{m}, \mathbf{y})$  delivers the message  $\mathbf{m}$  to the home-agent  $\mathbf{y}$ . The most common use of this operation is in a ruling for event  $\text{arrived}(\mathbf{x}, \mathbf{m}, \mathbf{y})$ , where operation  $\text{deliver}$  (with no arguments) simply delivers the arriving message to the home agent.

<sup>2</sup>Note that Moses regulates some additional types of events described in [9].

**The global set of rules  $\mathcal{R}$ .** The function of  $\mathcal{R}$  is to evaluate a ruling for any possible regulated-event that occurs at an agent with a given control-state. In our current model,  $\mathcal{R}$  is represented by a very simple Prolog-like program—or, if you will, a set of situation-action rules. When this “program”  $\mathcal{R}$  is presented with a goal  $e$ , representing a regulated event, and with a variable  $CS$ , representing the control-state of the home agent, it produces a list of primitive-operations representing the ruling of the law for this event.

### 3.2 Ticket Based Payment Policy—an Example

Consider the following token-based payment policy for the use a database accessed via a set of distributed servers:

Consider a distributed database that charges a fixed, and small, fee per query. Now, as is well known [4], e-cash certificates are not suitable for small payments, because each such certificate needs to be created by a bank, and each needs to be validated by the issuer, for fear of duplication. Therefore, let the payment be carried out according to the following policy:

A prospective client  $c$  can purchase from a designated ticket-seller an uncopyable  $q$ -valued ticket, where  $q$  is the number of queries this ticket is good for. A query submitted by  $c$  is processed only if  $c$  has a ticket with non-zero value, and as a result the value of this ticket is reduced by one.

This, to be called  $\mathcal{T}$ -policy, is formalized by law  $\mathcal{L}_{\mathcal{T}}$ , displayed in Figure 1. This law consists of six rules. Each rule in this figure is followed with a comment (in italic), which together with the following discussion, should be understandable even for a reader not familiar with Prolog. We now describe this law by explaining how various operations are regulated by it.

First, by Rule  $\mathcal{R}1$ , a client  $c$ , wishing to buy a ticket with value  $n$  sends to the ticket seller  $v$  the message `buyTicket(n,ec)`, where `ec` is a certificate for a sum of dollars in some form of electronic cash, like `Ecash` [11]. By Rule  $\mathcal{R}6$ , when such a message arrives at its destination it is delivered, only if the destination is an authorized ticket seller. The ticket seller may check the electronic certificate and if it is satisfied, it is expected to respond with a `giveTicket(n)` message (Rule  $\mathcal{R}2$ ), otherwise it is expected to return the electronic certificate. By Rule  $\mathcal{R}3$ , when the client receives a `giveTicket(n)` message a memo is delivered to the client, and the term `ticket(n)` which serves as a representation of the ticket prescribed by  $\mathcal{T}$ -policy, is added to its control state.

Next, by Rule  $\mathcal{R}4$ , a client  $c$  that has a ticket with positive value in his control state is allowed to present a query to any the servers of the database; the value of this ticket is automatically decreased by one unit. Finally, when a server responds to the query, the response is delivered to the client without further ado. (Rules  $\mathcal{R}5$  and  $\mathcal{R}6$ )

The important aspects of this law are that it insures that (1) only a ticket seller can inject tickets into the system, and (2) these tickets cannot be duplicated by clients, because they are kept by the controllers, and the law does not provide for their duplication. This is critical for the scalability of the mechanism, because a server may serve any query it receives *without having to check* the validity of the ticket. If servers are distributed, like in our case, the only solution we are aware of for preventing double spending is to have a centralized ticket checker, which tends to become a bottleneck [12].

Note that under this policy, the clients are at the mercy of both the ticket sellers and database servers. They send their payment with their order with no guarantee for either timely service or for the correct response. This problem can be alleviated by using the money-back policy discussed in Section 2 and implemented in Section 5.

### 3.3 The Distributed Enforcement Mechanism

The law for a given policy  $\mathcal{P} = \langle \mathcal{M}, \mathcal{G}, \mathcal{L} \rangle$ , is enforced in principle as follows: there is a *controller* associated with each member of group  $\mathcal{G}$ , logically placed between the agent and the communications medium, as is illustrated in Figure 2. All controllers have identical copies of the global set of rules  $\mathcal{R}$  of  $\mathcal{L}$ , and each controller maintains the control states of the agents under its jurisdiction. The steps taken when a member  $x$  wishes to send a  $\mathcal{P}$ -message  $m$  to a member  $y$  are:

1.  $x$  sends  $m$  to its assigned controller. The controller evaluates the ruling of the law  $\mathcal{L}$  for the event `send(x,m,y)` and it carries out this ruling. If part of the ruling is to forward a message  $m'$  to  $y'$ <sup>3</sup>, the controller sends  $m'$  to the controller assigned to  $y'$ .
2. when  $m'$  arrives to  $y'$ 's controller it generates an `arrived(x,m',y')` event. The ruling for this event is computed and carried out. The message  $m'$  is be sent to  $y'$  if so required by the ruling.

<sup>3</sup> $m'$  might differ from the original message  $m$ , and also the final destination  $y'$  might vary from the original destination  $y$ .

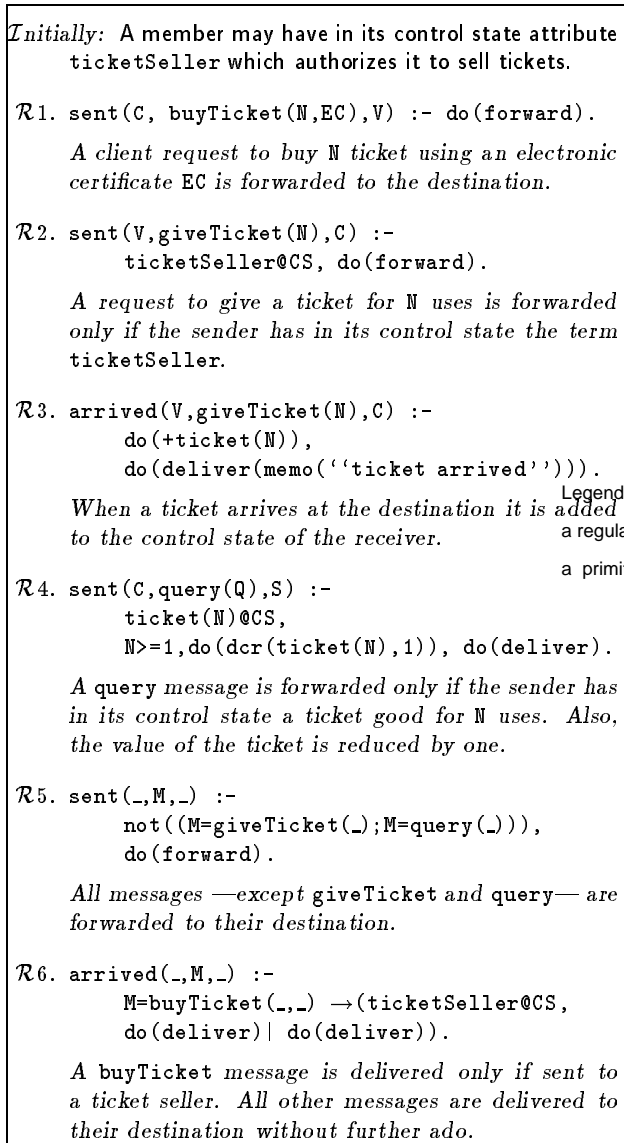


Figure 1: Law  $\mathcal{L}_T$  of Ticket Policy

The essential aspect of this architecture is that *all controllers have identical copies of the law*. It is in this sense that the law is said to be global to the group.

### 3.4 The Deployment of Policies, and their Scalability

A policy  $\mathcal{P}$  is introduced by creating a server that provides persistent storage for the law  $\mathcal{L}$  of this policy, including the control states of its members. This server is called the *secretary* of  $\mathcal{P}$ , to be denoted by  $\mathcal{S}_P$ .

For an agent  $x$  to be able to exchange  $\mathcal{P}$ -messages under a policy  $\mathcal{P}$ , it needs to send a connect message to  $\mathcal{S}_P$ , asking to be admitted either as one of the

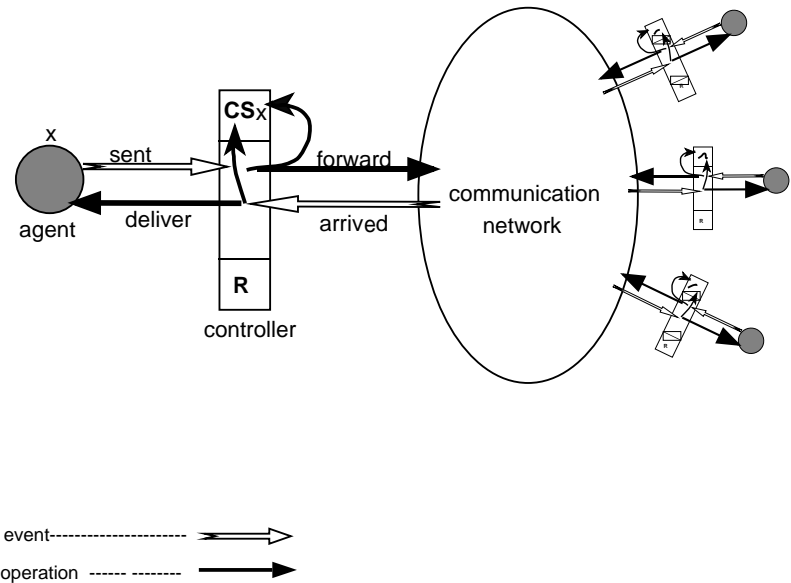


Figure 2: Enforcement of the Law

existing members of the group, which already has a name and a control state of its own, or as a new member, requesting a generic control state.  $\mathcal{S}_P$  is likely to require authentication, particularly for an attempt to connect to an existing member, using some standard cryptographic techniques. If the secretary agrees to admit  $x$ , it would assign him a controller  $\mathcal{C}_x$ , selected from a set of *trusted controllers*, as will be explained later. By itself, this controller knows nothing about our policy, but  $\mathcal{S}_P$  will feed  $\mathcal{C}_x$  with the set of rules  $\mathcal{R}$  of the law  $\mathcal{L}$ , and with the control state assigned to  $x$ .

Once this connection is made, the interaction of  $x$  with the other members of policy  $\mathcal{P}$  is handled by  $\mathcal{C}_x$ , and does not directly involve the secretary  $\mathcal{S}_P$ . *This makes regulated interaction quite scalable.*

Several clarifying points are in order here. First, nobody is ever forced to join a given policy, such as our  $\mathcal{MB}$ . But if a vendor insists on accepting orders only as  $\mathcal{MB}$ -messages, his clients have no choice but use  $\mathcal{MB}$ -messages for making such orders. Conversely, if the clients insist on having a money-back guarantee, then the vendor would have no choice but providing such a service. In general, if a policy is deemed useful by one side or the other, then it is likely to be adopted by all participating agents.

Second, note that nothing prevents a member of one policy-group, such as  $\mathcal{MB}$ , to join another one, or to use unregulated messages for part of its communication. Therefore, a given agent may participate in any number of such policy-groups.

Finally, we note that a policy does not have to be

supported by a single secretary. It is possible, in principle, for a policy to have several secretaries, each maintaining a subgroup of  $\mathcal{G}$ , perhaps in a different geographic region of the Internet.

## 4 Security and Trust

The security of the proposed mechanism is based on the following assumptions: (1) messages are securely transmitted over the network; (2)  $\mathcal{P}$ -messages are sent and received only via correctly implemented controllers, interpreting law  $\mathcal{L}_{\mathcal{P}}$ ; and (3) the secretary  $\mathcal{S}_{\mathcal{P}}$  can be trusted by everybody who depends on it. To ensure the first of these conditions, each controller, and potentially each agent in  $\mathcal{G}_{\mathcal{P}}$ , has a pair of (RSA) keys: a public key known to a trusted authority and a secret key known only by itself. If the messages sent across the network are digitally signed, their authenticity is guaranteed as long as the private key is not disclosed.

For the sake of the second condition above we postulate the existence of a “public utility” of *trusted controllers* which we call a *controller-service*. It is an organization that maintains a set of trusted hosts that function as controllers. The controller-service also maintains one or more name servers, called *C-servers*, that act as an authentication agents for controllers. These trusted controllers are generic, and can interpret any law provided to them by a secretary  $\mathcal{S}_{\mathcal{P}}$  of a given policy  $\mathcal{P}$ . Besides interpreting laws, the controllers are built to communicate with secretaries, with members of a given policy-group, and with each other. Each such communication requires its own specific protocol. The controller-controller authentication protocol is particularly important; it needs to ensure each controller that its peer is a valid controller (i.e., authenticated as such by a C-server) and that it operates under the same policy.<sup>4</sup>

To be effective, such controllers need to be well dispersed geographically, so that it would be possible to find a controller reasonably close to prospective users. And there needs to be a trusted server that will identify a suitable controller when requested to do so by a secretary  $\mathcal{S}_{\mathcal{P}}$ . We call it a C-server (C for “controller”). We believe that such a utility can be established by a large, trustworthy financial institution – like Visa, MasterCard or major banks– or by a large Internet provider.

The trustworthiness of controllers is due to several factors, besides the degree to which one is willing to

<sup>4</sup>The protocol we have devised, which cannot be described here due to space limitations, includes a hash of the law of  $\mathcal{P}$  in each message, and uses digital signature of the messages themselves, in a standard manner.

trust the controller service. First, it is difficult, although not impossible, to commit fraud by subverting controllers, because they are not designed for any specific application. Second, controllers can be tested anytime by loading them with a policy designed specifically for testing.

Finally, regarding condition 3 above, we point out that a secretary  $\mathcal{S}_{\mathcal{P}}$  may have to be trusted only by some members of its policy-group, because it cannot do much to defraud the others. First, because  $\mathcal{S}_{\mathcal{P}}$  cannot do any harm by supplying altered rules to a controller, because such a controller would be unable to communicate with other controllers that do operate under the correct law—this is due to our controller-controller protocol. Besides, a suspicious client, who should be familiar with the law of the policy under which it operates, can always ask the controller assigned to it to show it its current law.

Thus, the only way for a secretary to cause serious harm is by associating incorrect control state to an agent, thus providing him with undeserved powers under the law of the policy in question. So, in the case of policies where members are “equal” under the law, in the sense, that all members start with the same control state, the secretary need not to be trusted by any member. A large collection of policies falls into this category, including the policies presented by Ketchpel and Garcia-Molina [5], and our money-back policy presented in 5. There are cases, however, where some members have special rights under a policy because of their control state. If this is the case, the secretary has to be trusted, at least by agents that might be affected by such special rights. One way to deal with this issue is to offer liability insurance to agents operating under the policy in question, in a manner described in [7].

## 5 Establishing the Money Back Policy

We introduce here law  $\mathcal{L}_{\mathcal{MB}}$ , displayed in Figure 3, which implements the money back policy  $\mathcal{MB}$  described in Section 2. We now describe this law by explaining how various purchase-related operations are regulated by it.

First, consider a client  $c$  who wants to purchase from vendor  $v$  a merchandise defined by  $\mathit{specs}$ , for which he is willing to pay 9 dollars, requesting a grace period of 60 minutes. This is accomplished by  $c$  sending a  $\mathit{order}(\mathit{specs}, 9, 60)$  message to  $v$ . Due to Rule  $\mathcal{R}1$  of  $\mathcal{L}_{\mathcal{MB}}$ , this message would be forwarded to  $v$  *only* if  $c$  sends a valid electronic certificate for \$9. Note that the electronic certificate is not sent to the vendor, but kept instead in the client’s control state by the term  $\mathit{escrow}(v, \mathit{specs}, 9, \mathit{ec}, 60)$ , thus effectively placing the certificate in escrow. By Rule  $\mathcal{R}2$ , when

<p><i>Initially:</i> A member may have in its control state an attribute vendor which characterizes the role it plays.</p> <p><math>\mathcal{R}1</math>. <code>sent(C,order(Specs,Fee,EC,GP),V) :-</code>  <code>do(+escrow(V,Specs,Fee,EC,GP),</code>  <code>checkMoney(Fee,EC),</code>  <code>do(forward(C,order(Specs,Fee,GP),V)).</code></p> <p><i>A message order(Specs,Fee,GP) is forwarded to the destination only if the sender has a valid electronic cash certificate EC for sum Fee.</i></p> <p><math>\mathcal{R}2</math>. <code>arrived(C,M,V) :- vendor@CS,</code>  <code>(M=order(.,.,.)   M=refund(.)  </code>  <code>M=payment(.)), do(deliver).</code></p> <p><math>\mathcal{R}3</math>. <code>sent(V,M,C) :- vendor@CS,</code>  <code>(M=supply(.,.)   M=paymentRequest(.))</code>  <code>do(forward).</code></p> <p><i>Messages arriving at a vendor, or sent by a vendor are delivered, respectively forwarded without further ado.</i></p> <p><math>\mathcal{R}4</math>. <code>arrived(V,supply(Specs,M),C) :-</code>  <code>clock(T),do(+supplied(V,Specs,T)),</code>  <code>do(deliver).</code></p> <p><i>If the vendor supplied the order, the merchandise is delivered to the client.</i></p> <p><math>\mathcal{R}5</math>. <code>sent(C,refund(Specs),V)</code>  <code>:- escrow(V,Specs,Fee,EC,GP)@CS,clock(T),</code>  <code>supplied(V,Specs,T')@CS <math>\rightarrow</math> T &lt;= T' + GP</code>  <code>  true,</code>  <code>do(deliver(V,refunded(Fee,EC),C)),</code>  <code>do(-escrow(V,Specs,Fee,EC,GP)),</code>  <code>do(forward).</code></p> <p><i>If a client asks for a refund either before the merchandise was delivered, or while the grace period is in effect then the electronic certificate with which the client paid is delivered to him and the correspondent escrow term is removed from the control state.</i></p> <p><math>\mathcal{R}6</math>. <code>arrived(V,paymentRequest(Specs),C)</code>  <code>:- escrow(V,Specs,Fee,EC,GP),clock(T),</code>  <code>supplied(V,Specs,T'), T &gt; T' + GP,</code>  <code>do(forward(C,payment(Specs,EC),V)).</code></p> <p><i>If the grace period ended a payment request is honored by forwarding the electronic certificate to the vendor.</i></p>
--

Figure 3: Law  $\mathcal{L}_{MB}$ —of Money Back Policy

this message arrives at  $v$ , it will be delivered.

The vendor  $v$  may respond to the order he received from  $c$ , with a `supply(specs,m)` message, where  $m$  is a certificate of some sort for the ordered merchandise (Rule  $\mathcal{R}3$ ). By Rule  $\mathcal{R}4$ , the arrival of such a message at  $c$  causes the following operations to be carried out: (1) the merchandise is delivered to the client; and (2) the time  $t$  when the merchandise was supplied is recorded by a term `supplied(v,specs,t)`, so that the grace period can be observed.

The client may ask for a refund, by sending a `refund(specs)` to the vendor  $v$ . If the vendor has not yet supplied the merchandise or the grace period has not ended, the certificate `ec` with which the client paid for the merchandise is removed from escrow, by removing the corresponding `escrow` term, and delivering the certificate to the client. The certificate can now be used in a new transaction.

Finally, to actually get its money, a vendor can ask for payment by the means of a `paymentRequest(specs)` message. When the message arrives at the client, if the grace period is over, it causes the removal from escrow of the electronic certificate that served as payment, and its sending to the vendor (Rule  $\mathcal{R}6$ ).

**Discussion** In this example, the merchandise was delivered to the client as soon as it was supplied by the vendor. If the merchandise is a certificate of some sort – like a flight ticket – the vendor incurs no losses if a refund is granted because it can always void the certificate. If the merchandise is of a different nature, like for example a digital document, the vendor may require that the actual merchandise is supplied to the client, only after the grace period ended. In this case, the law can be modified so that the merchandise is kept in escrow on the vendor’s side and delivered only when the payment is received.

## 6 Related Work

Generally speaking, this paper is based on, but largely orthogonal to, the large volume of work on cryptographic techniques and on payment mechanisms across the net. It should be enough to mention here the following three payment methods: electronic currency [3, 8, 11], debit/credit instruments [10, 12], and secure credit card. Although our examples used electronic cash, any other payment methods could have been used.

What is more relevant to this paper is that several researchers recently addressed the need to implement higher level e-commerce policies; a careful analysis of

this need has been given by Waidner [15]. All the proposed implementations we have seen for such policies employ *trusted proxies*—like our controllers—to mediate between the principals involved in a given commercial activity, making sure that the intended policy is not violated. These proposals differ from each other, and from the present proposal, in the manner they use these proxies, and in the nature and generality of their concept of a policy.

Ketchpel and Garcia-Molina [5] studied the transactions that occur between a customer who buys items from different vendors through brokers. The integrity of such transactions is ensured by trusted agents placed between every two principals, which operate according to rules generating by a technique called *graph sequencing*. This is an effective technique, but it is limited to a certain class of activities, and is concerned with a certain kind of issues. For example, their technique cannot prevent duplication of certificates, nor can they provide for auditing.

Su and Manchala [13] also deal with a specific class of transactions, using trusted agents to mediate between a client and a vendor. An important difference between this work and ours is that their trusted agents are not generic. Each controller must be programmed to carry out a specific protocol, which may be different from agent to agent. The deployment and management of such heterogeneous agents, and the maintenance of their trustworthiness, may be difficult. Also, Su and Tygar [14] used trusted agents to solve certain fault tolerance problems.

Examples of other electronic commerce applications that use non-generic trusted agents are Millicent [4] and NetBill [12]. In Millicent, client pay vendor with vendor's scrip obtained from a broker. The brokers are entities trusted to perform all exchanges correctly. In the latter protocol, Sirbu and Tygar propose a business model where merchants and clients open accounts with a NetBill server, which is the entity trusted to commit financial transactions.

Finally, Blaze, Feigenbaum and Lacy [2] built a toolkit called PolicyMaker which can interpret security policies. An agent (usually a vendor) receiving a request gives it for evaluation to PolicyMaker together with its specific policy, and the requester's credentials. On this basis the request can be found to be valid, invalid or trust can be deferred to third parties. One of the main differences between this work and ours is that PolicyMaker provides no enforcement. In particular, after asking PolicyMaker for its ruling one can proceed by ignoring it.

## 7 Conclusion

What makes e-commerce riskier, and thus more difficult, than conventional commerce is that its participants generally do not trust each other, because, as has been pointed out in [5], *“since the principals [in an electronic transaction] may have no previous working relationship, nor any expectation of continued interaction in the future, they may be distrustful of each other”*.

What we have proposed here is to engender trust between the participants in a given e-commerce activity by imposing a suitable law over the interaction between them. The difference between unregulated e-commerce and one that is regulated by an enforced law, is analogous to the difference between doing business in an anarchic, lawless country, and doing it within a lawful society.

The benefits of regulated e-commerce can be demonstrated with our  $\mathcal{MB}$  policy. The confidence that a potential client has, under this policy, of getting his money back, if he asks for it in time, is likely to embolden him to make orders he might otherwise be reluctant to make. This is true, in particular, in the following two cases: First, a potential client  $c$  may be reluctant to make an order from a server whose response time is uncertain, fearing that his money will be committed for a long period of time before getting the ordered merchandise—if he will get it at all. Under  $\mathcal{MB}$ , on the other hand,  $c$  can be sanguine about making orders, because he can always get his money back if not serviced quickly enough.

The second benefit to be considered is a simple solution for the following well known problem of coordination in e-commerce (this problem has been addressed, in particular, in [5, 1]). Suppose that  $c$  wants to purchase a set of items from different vendors, but he wants these items *only* if he can get them all. (These items might be tickets for a series of flights, planned for a single trip, or a set of related documents). The problem here is to avoid the danger of having  $c$  stuck with some of the items he wants, for which he has already paid, not being able to get the rest of them, for some reason. The conventional solution to this problem is to have this transaction carried out by some kind of coordinator, who must be trusted both by  $c$  and by all the vendors involved. But no such coordinator is required under  $\mathcal{MB}$ . If  $c$  purchases all these items with a sufficient period of grace, then he can get refunds for the items he purchased if he decides that he cannot get all the items he needs. A little bit of trust can sometimes go a long way.

It should be pointed out that there is a price to

be paid for this gain of trust, which consists of: (a) the run-time cost of having each regulated message go through a pair of controllers; and, (b) the effort of maintaining a collection of identical trusted controllers throughout the Internet. We believe that the first of these costs is quite acceptable, but it is still an open question whether it is practical to deploy such controllers on a truly large scale.

## References

- [1] J.-M. Andreoli, F. Pacull, and R. Pareschi. XPECT: A framework for electronic commerce. *IEEE Internet Computing*, pages 40–48, July–August 1997.
- [2] M. Blaze, J. Feigenbaum, and J. Lacy. Decentralized trust management. In *Proceedings of the IEEE Symposium on Security and Privacy*, May 1996.
- [3] D. Chaum. Transaction systems to make big brother obsolete. In *Communication of the ACM*, October 1985.
- [4] S. Glassman, M. Manasse, M. Abadi, P. Gauthier, and P. Sobalvarro. The Millicent protocol for inexpensive electronic commerce. In *Fourth International World Wide Web Conference Proceedings*, pages 603–618, December 1995.
- [5] S. Ketchpel and H. Garcia-Molina. Making trust explicit in distributed commerce transactions. In *Proceedings of the International Conference on Distributed Computing Systems*, pages 270–281, 1996.
- [6] S. Ketchpel, H. Garcia-Molina, and A. Paepcke. Shopping models: A flexible architecture for information commerce. In *Digital Libraries*, 1997.
- [7] C. Lai, G. Medvinsky, and C. Neuman B. Endorsements, licensing, and insurance for distributed system services. In *Proceedings of the Second ACM Conference on Computer and Communication Security*, November 1994.
- [8] G. Medvinsky and C. Neuman. Netcash: A design for practical electronic currency on the internet. In *Proceedings of the 1st ACM Conference on Computer and Communication Security*, 1993.
- [9] N.H. Minsky and V. Ungureanu. Regulated coordination in open distributed systems. In David Garlan and Daniel Le Metayer, editors, *Proc. of Coordination'97: Second International Conference on Coordination Models and Languages; LNCS 1282*, pages 81–98, September 1997.
- [10] C. Neuman. Proxy-based authorization and accounting for distributed systems. In *Proceedings of the 13th International Conference on Distributed Computing Systems*, 1993.
- [11] P. Panurach. Money in electronic commerce: Digital cash, electronic fund transfer and ecash. *Communications of the ACM*, 39(6), June 1996.
- [12] M. Sirbu and J.D. Tygar. Netbill: An Internet commerce system. In *IEEE COMPCON*, March 1995.
- [13] J. Su and D. Manchala. Building trust for distributed commerce transactions. In *17th IEEE International Conference on Distributed Computing Systems (ICDCS)*, May 1997.
- [14] J. Su and J.D. Tygar. Building blocks for atomicity in electronic commerce. In *Proceedings of USENIX Security Symposium*, 1996.
- [15] M. Waidner. Development of a secure electronic marketplace for Europe. In *Proceedings of Esorics*, September 1996.