

Decentralized Governance of Distributed Systems via Interaction Control

Naftaly H. Minsky

Rutgers University

Abstract. This paper introduces an abstract reference model, called *interaction control* (IC), for the governance of large and heterogeneous distributed systems. This model goes well beyond conventional access control, along a number of dimensions. In particular, the IC model has the following characteristics: (1) it is inherently decentralized, and thus scalable even for a wide range of stateful policies; (2) it is very general, and not biased toward any particular type of policies; thus providing a significant realization of the age-old principle of *separation of policy from mechanism*; and (3) it enables flexible, composition-free, interoperability between different policies.

The IC model, which is an abstraction of a mechanism called law-governed interaction (LGI), has been designed as a minimalist reference model that can be reified into a whole family of potential control mechanisms that may support different types of communication, with different performance requirements and for different application domains.

1 Introduction

The economy and security of modern society depends increasingly on the software that supports complex institutions, such as banking systems, health care organizations, government agencies, and commercial enterprises. Such software is undergoing a relentless transition from monolithic systems, constructed according to a single overall design, into heterogeneous, loosely coupled, distributed systems, whose component parts may be written in different languages, may run on different platforms, and may be designed, constructed, and even maintained under different administration domains. Prominent examples of this trend include: *heterogeneous multi-agent systems*, and systems based on the concept of *service oriented architecture* (SOA).

Unfortunately, the semi-anarchical nature of such systems engenders serious obstacles to their dependability and security. These obstacles include, but are not limited to, the following: (1) The difficulty to protect a heterogeneous system from attacks by unruly or malicious *actors*, from within the system or from the outside—where by an “actor” we mean in this paper either an autonomous software component, or a person (or even some physical device) operating via some software/user interface. And (2) the difficulty to facilitate safe and harmonious

coordination between the system’s disparate actors, which may have little or no knowledge of each other’s behavior.

There is a broad consensus that some of these difficulties can be alleviated by suitable control of the interaction between the disparate actors of a system, and without assuming any knowledge of, or control over, the structure or internal behavior of the interacting actors themselves. The issue addressed in this paper is the nature of control mechanism that is suitable for this purpose.

The conventional approach to the control of distributed systems is based on the concept of access control (AC), rooted in the *matrix model* formulated about forty years ago [18]. But we shall argue in this paper that this model, which has been designed for use within a single computer host, is unsuitable for the control of distributed systems; and that this goes for the contemporary AC mechanisms based on this model. It should be pointed out, however, that we are not referring here to what may be called *server-centric*¹ AC, which deals with a single server that has a policy regarding the services it intends to provide to its distributed clients. Rather, in this paper we address only *system-centric*² AC mechanisms that aim to regulate the flow of messages in an entire system, rather than, or in addition to, leaving it up to individual servers to decide on their own AC policies.

We will mention here two of the limitations of the conventional concept of AC, when applied to distributed systems. First, despite the distributed nature of the system being controlled, the control policies under AC are enforced via a virtually *centralized reference monitor*³. It is obvious that such centralization increases the vulnerability of the control mechanism to attacks, and has other drawbacks to be pointed out in due course.

Second, like the matrix model on which it is based, AC does not, generally⁴, support *stateful policies*, which are sensitive to the history of interaction. Indeed, AC has been designed *to permit access only to those that have the right for it*, but without any concern about the the dynamic behavior of those that get such access. This is analogous to having traffic laws that require only that every driver has a driving license, without any concern about how he drives, and how he reacts to inputs from his dynamically changing environment—which is, of course, critical for the security of driving. Moreover, any attempt to support stateful policies under AC is bound to be unscalable, due to the centralized nature of AC, as we shall see in Section 4.1.

Given these drawbacks of conventional AC it behooves us to reconsider the very nature of control over distributed systems. We do it in this paper in the following manner. We start in Section 2 by proposing a set of principles that such control should satisfy. We call them the the *interaction-control* (IC) principles, as we insist that it is necessary to control the dynamic process of interaction

¹ This type of AC is exemplified by *trust management*, such as under Keynote [10], and by the BPEL policy language of SOA.

² Exemplified by XACML [14]

³ “Virtually,” because the reference monitors are often replicated.

⁴ with some exceptions to be noted later.

between actors, and not—as under AC—be limited to blocking access to those that do not have the *rights* to it. In Section 3 we introduce a reference model of a control mechanism that satisfies these principles; it is called the *IC model*. This model is an abstraction of the implemented *law governed interaction* (LGI) mechanism [22, 25, 24]—featuring only those aspects of LGI that are essential for the support of our principles, and ignoring other important aspects of this mechanism, as well as a host of practical details necessary for a real mechanism such as LGI.

There are two reasons for introducing this abstract model. First, this model is to serve as a basis for a whole family of concrete *IC mechanisms* that would satisfy the IC principles, but may be designed for different types of communication, different performance requirements, and different application domains; and which may extend the minimalist IC model in various ways. The second reason is that due to the abstract nature and the relative simplicity of the IC model, it lends itself to analysis. We exploit this fact in Section 4, where we explain how the various IC principles are made to be satisfied by the IC model; and we expose the underlying differences between access control, and the mechanisms based on the IC principles. The paper concludes in Section 5.

2 The Principles Underlying Interaction Control

We propose here five principles that we require an interaction control mechanism to satisfy in order to support an effective governance of distributed systems. The statement of these principles is followed with some clarifying and motivating comments. We elaborate further on these principles, and on their broader implications, throughout this paper.

1. **Sensitivity to the history of interaction:** Interaction control needs to support policies that are sensitive to the history of interaction between the actors operating in the system. (Such a mechanism is often called *stateful*, because being aware of the history of interaction requires the maintenance of a state.)
2. **Decentralization:** An IC mechanism should be decentralized, in the following sense. The enforcement of system policy over the interactive activities of a given actor, should be carried out *at that actor*, independently of the coincidental state and behavior of any other system actor.
3. **Global sway:** Despite its decentralized nature, an IC mechanism should be able to have global sway over the system under its jurisdiction.
4. **Flexible interoperability:** An IC mechanism should provide flexible means for actors operating under different policies to interoperate, without having to compose their respective policies into a single one.
5. **Separation of policy from mechanism:** The IC mechanism should not be biased towards any particular type of policies.

Discussion: We already motivated the first two principles above. The global sway required by Principle 3 is obviously important, and it seems to contra-

dict Principle 2. But as we shall see later, these two principles can be satisfied together. The following are some comments about the last two principles.

Interoperability—the subject of Principle 4— is critical due to the growing need for actors operating under disparate policies to interact. This is the case, for B2B commerce, and for various kind of federation of organizations. Interoperability is also required in very complex system, which are governed by several, sometimes many, policies. The need to be able to interoperate without having to compose the policies of the interacting parties into a single policy is due to the fact that composition is computationally hard [20] and tends to be very inflexible, as argued in [2].

Finally, Principle 5, which requires *separation of policy from mechanism* is an age-old principle formulated more than thirty year ago [32] for access control, but never fully satisfied. It means, essentially that a control mechanism should not be designed specifically for any particular type of policies—such as particular ways for doing delegation and revocation of rights. But it should instead be able to support a wide range of policy types, in a uniform manner. This principle is becoming increasingly important because modern complex systems tend to require a multitude of diverse policies, regarding various types of system activities, and various system divisions. Using different mechanisms to implement such policies would be very hard, and it would preclude interoperation between actors subject to different types of policies.

3 The Interaction Control (IC) Model

We introduce here an abstract model of mechanisms for the control of the interactions (via message exchange) between distributed actors. This model, which is based on the IC principles, is oblivious of the internal structure and behavior of the interacting actors, and independent of the languages in which these actors are programmed, and of the type of host on which they run. In other words, the actors whose interactions are to be regulated are viewed as black boxes that send and receive messages. (This enables the IC model to deal with program modules and people in completely uniform manner.)

The IC model is abstract in the following respects. First, it does not specify the type of message passing being used for inter-actor interactions; which may, in particular, be asynchronous, synchronous (via RPC of some kind), or both. Second, this model makes no assumptions about the reliability and latency of the message passing, or about the underlying fabric of the communication network—which may, in particular, be the Internet, or some form of ad hoc wireless communication. Third, this model leaves many details unspecified, which makes it very lean.

The most basic elements of this model are: (a) the concept of a *interaction law*, (or, simply a *law*), which is analogous to, although fundamentally different than, the concept of policy under AC; (b) the concept of an *\mathcal{L} -agent*, which, broadly speaking, is an actor engaged in sending and receiving messages, subject to a specific law \mathcal{L} (we often use the term *IC-agents*, or simply an *agent*, when the law under which it operates is assumed to be known, or if it is of no immediate

concern); and (c) the concept of *law-based trust*, or *L-trust* for short, induced by an IC mechanism.

This section is organized as follows. After introducing the concepts of law and of agent in the following two sections, we discuss in Section 3.3 the dual mediation mechanism between any pair of IC-agents implied by this model. We then illustrate these concepts in Section 3.4 via an example. In Section 3.5 we introduce the concept of L-trust, and show how it is used for the enforcement of IC laws. Finally, in Section 3.6 we discuss possible techniques for establishing L-trust.

3.1 The Concept of Law

The function of a law \mathcal{L} under this model is to decide what should be done in response to the occurrence of certain events at any agent x that operates under this law (i.e., at any \mathcal{L} -agent). This decision is called the *ruling* of the law for the event e that triggered it. More specifically, a law \mathcal{L} that governs an agent x is defined in terms of three elements: (1) a set E of *regulated events* (or, simply, “events”) that may occur at any agent; (2) the *control-state* CS_x associated with each agent x , which is distinct from the internal state of x (which is invisible to the law); and (3) and a set O_i of *control operations*, that can be mandated by a law. These elements are defined abstractly below—and they need to be reified, and can be extended, by any IC mechanism based on this model.

The Set E of Regulated Events: This is the set of events that may occur at any agent and whose disposition is subject to the law under which this agent operates. And every event is associated with the local time of its occurrence. This set contains at least the following three types of events, described here with little, if any, detail of their structures: (1) The **arrived** event, which represents the arrival of a message at a given agent, and contains an identifier of the law under which this messages has been sent. (2) The **sent** event, which represents an attempt by the actor of a given agent to send a message (the distinction between an agent and its actor will be clarified in Section 3.2). And (3) the **birth** event, which signals the moment of formation of the agent in question, and is the very first regulated event in its lifetime. (Among the important additional events under the LGI mechanism—which are not included in this model—are *the coming due of an obligation*, and *the occurrence of an exception*.)

The Control-State of an Agent: The control-state CS_x of an agent x is the state maintained by x —it is distinct from the internal state of the actor of this agent, of which this model is oblivious. The control-state, or simply “state,” is an unbounded bag of terms whose structure is not specified by this model (As we shall see, this state, which is initially empty, can change dynamically in response to the various events that occur at it, subject to the law under which this agent operates.)

The Set O of Control Operations: These are the operations that can be included in the rulings of laws. They contain the following operations, which affect the exchange of messages between the home agent and its peers: (a) the **forward** operation, which forwards a message to another agent; and (b) the

`deliver` operation, which delivers a message to the actor of the home agent—in effect, enabling this actor to access this message.

We are now in a position to give a more formal definition of the concept of law, as follows.

Definition 1 (law). *Given the sets E and O , as defined above, and a set S of all possible control-states of an agent, a law \mathcal{L} is a mapping:*

$$\mathcal{L} : E \times S \rightarrow S \times (O)^* \tag{1}$$

In other words, for any given (e, cs) pair, the law mandates a new state (which may imply an arbitrary state change), as well as a (possibly empty) sequence of control operations that are to be carried out in response to the occurrence of event e in the presence of state cs . Several aspects of this concept of law are worth pointing out here.

1: There is an interplay between the fixed law that governs a given agent, and its dynamically changing state: on one hand, the ruling of the law may depend on the state, and on the other hand the content and evolution of the state is regulated by the law. This interplay is an important characteristic of the IC model, which, as we shall see in Section 4, is instrumental in satisfying Principles 1 and 2.

2: *A law is local* at every agent x it governs, in that the ruling of the law is defined over the local events and state at x , and it can mandate only operations that can be carried out locally at x . Thus the law that governs an agent x can be complied with, or enforced, locally at x , without any access to other agents, and without any knowledge of the coincidental state of others—as required by Principle 2. It is this locality which makes the IC-model decentralized. Moreover, as we shall see in Section 4.5, the locality of the law facilitates flexible interoperability between agents subject to different laws, as required by Principle 4.

3: A law cannot be inconsistent⁵; that is, its ruling is never ambiguous. Indeed, by its definition, a law maps any given $(event, state)$ pair to a *single* ruling. This is worth mentioning, in particular, because many conventional policy mechanisms admit policies with inconsistencies, which need to be resolved, somehow. The main reasons for inconsistencies under these mechanisms, is the wish to admit independent contributions to a single policy by different stakeholders. But this need can be satisfied under IC, without introducing inconsistencies, by extending the IC model with the concept of LGI-like *law-hierarchy*, (cf. [2]).

4: Finally, note that this abstract definition of the concept of law does not specify a language for writing laws (which we call a law-language). This is for several reasons. First, despite the pragmatic importance of choosing an appropriate law-language, this choice has no impact on the semantics of the model itself, as long as the chosen language is sufficiently powerful to specify all possible

⁵ This does not mean, of course, that the a law cannot be incorrect, in the sense that it does not function as its writer intended.

mappings of the form defined above. Second, by not specifying a law-language one provides the freedom to employ different law-languages for different applications domains, possibly under the same mechanism. (Indeed, the implemented LGI mechanism employs two different law-languages, one based on the logic-programming language Prolog, and the other based on Java.) Finally, leaving the language out of this model simplifies it, and facilitates the analysis of its underlying nature. Note, however, that for the sake of the examples used in this paper we will introduce an informal pseudo code for such a language.

3.2 The Concept of an \mathcal{L} -agent

We have broadly defined an \mathcal{L} -agent as an actor engaged in the sending and receipt of messages, in compliance with law \mathcal{L} . Here we elaborate on this definition. First, by introducing an abstract model of an \mathcal{L} -agent, second by defining the meaning of compliance with a given law, and concluding with a brief discussion of cross-interaction (also called “interoperation”) between agents operating under different laws.

A Model of an \mathcal{L} -Agents We define an \mathcal{L} -agent x to be a pair $x = \langle c_x, T_x^{\mathcal{L}} \rangle$, where c_x is an *actor*, and $T_x^{\mathcal{L}}$ is a device that mediates the interactions of c_x with others, in compliance with law \mathcal{L} . The mediator $T_x^{\mathcal{L}}$ is called the *private controller* (or simply the controller) of agent x , and the actor c_x , whose internal structure is irrelevant to this model, is said to *animate* agent x . (Figure 1 depict two such agents, operating under possibly different laws, and interacting with each other.)

Note that this abstract model does not specify the “geographical” relationship between an actor and the controller that mediates its interaction with others. Indeed, as we shall see in Section 3.6, different realizations of this model may have the controller either running on a different host than the actor it serves, placed anywhere on the Internet; or running on the host of the actor; or even be part of the actor itself. Also, we point out that the controller $T_x^{\mathcal{L}}$ should not be viewed as a *wrapper* of actor c_x . In particular, because we do not rule out the possibility that an actor animates, concurrently, several agents, possibly under different laws; and may, at the same time, communicate in a manner not regulated under IC.

Now, it is the controller $T_x^{\mathcal{L}}$ which is the locus of the *regulated events* at agent x . In particular, an **arrived** event occurs at controller $T_x^{\mathcal{L}}$ when a message addressed to x arrives at $T_x^{\mathcal{L}}$. Similarly, a *sent* event occurs at $T_x^{\mathcal{L}}$ when a message sent by its actor c_x arrives at it, on its way to its target. It is also the controller $T_x^{\mathcal{L}}$ that maintains the *control-state* of agent x , and carries out the ruling of the \mathcal{L} .

Structurally, a controller can be described as a triple $T_x^{\mathcal{L}} = \langle \mathcal{L}, I, CS_x \rangle$, where \mathcal{L} is the law under which this particular controller operates; CS_x is the *control-state* of agent x ; and I is a generic mechanism that complies with any given law \mathcal{L} by interpreting this law, and carrying out its rulings, in the manner defined below.

Compliance with a Law: We say that an agent x complies with a given law \mathcal{L} if the following two conditions are satisfied: (1) The ruling of law is carried out⁶, *atomically*, in response to the occurrence of every regulated event at x , and in the order of the occurrence of these events; while events that occur simultaneously are handled sequentially, in arbitrary order. And (2) control operations and state changes are *not* carried out at x unless mandated by its law—thus, in particular, no messages would be forwarded from x to anybody, without it being mandated by the law.

\mathcal{L} -Messages, and Cross-Interaction: By definition, messages sent by an \mathcal{L} -agent comply with law \mathcal{L} —they are thus called \mathcal{L} -messages. But an \mathcal{L} -agent may be permitted, by its law \mathcal{L} , to receive \mathcal{L}' -messages—that is, messages sent by an \mathcal{L}' -agent—where law \mathcal{L}' is different from \mathcal{L} . This amounts to *cross-interaction* between agents operating under different laws. Moreover, an \mathcal{L} -agent may be permitted, by its law \mathcal{L} , to receive even *unregulated messages*, that is, messages sent by an actor not regulated by an IC mechanism⁷.




Fig. 1. A pair of interacting agents, operating under possible different laws

3.3 The Dual Mediation of Communication Under the IC model

One of the significant aspects of the IC model is that it involves dual mediation of every exchange of messages between IC-agents: one on the side of the sender of a message, and one on the side of its receiver. Specifically, the passage of a message from an actor c_x of an \mathcal{L} -agent x to an actor c_y of an \mathcal{L}' -agent y , must be mediated first by the controller $T_x^{\mathcal{L}}$ associated with c_x , and then by the controller $T_y^{\mathcal{L}'}$, associated with c_y , as is illustrated in Figure 1. This is a direct consequence of the locality of IC laws, which requires both the sender and receiver to individually comply with the law under which each of them operates. This dual mediation is in contrast with the conventional AC mechanisms, which use a single *reference monitor* to mediate the exchange of messages. Such a reference monitor is usually placed at the server side, or is used as a central *policy decision point* (PDP) for an entire system, as under XACML [14].

⁶ If the ruling of the law for a given event is empty, then nothing extra will be done in response to this event; in other words, this event is effectively ignored by the law.

⁷ For simplicity, we mostly ignore in this paper the latter possibility, which is quite important in practice.

The dual mediation under IC has several important implications, not the least of which is that it facilitates interoperability by providing flexible control over cross-interaction between agents operating under different laws, as is further discussed in Section 4.5. Moreover, as has been shown in [25], the dual control turns out to be more efficient than centralized control, in many circumstances. A simple illustration of the nature of dual mediation, and some of its consequences, is provided by the following example law.

3.4 Budgetary Control Over Messaging—an Example:

We introduce here a rather synthetic law that budgets messaging activities. More specifically, this law, called \mathcal{BC} , for “budgetary control,” limits any actor that operates under it to sending at most 1000 messages, and to receiving no more than 2000 messages.

This, and other example laws in this paper, are written here in an informal pseudo-code, which is broadly modeled after the logic-based law-language used by the LGI mechanism [24]. We first introduce this pseudo code, and then law \mathcal{BC} written in it.

An Informal Law-Language: A law written in this language consists of a sequence of *event-condition-action* rules of the following form:

UPON e IF c DO $[o]$,

where e is a regulated event (or a pattern that could match several such events); c is an optional condition defined over the event e , and over the state (CS) of the *home agent*—i.e., the agent served by the controller in question; and $[o]$ is the sequence of control-operations and state-changing operations that constitute the ruling of the law (i.e., the “action” part of the rule).

Whenever some event e' occurs at some agent x , the ruling of its law is computed by the controller of this agent in the following way: The sequences of rules that constitute the law are evaluated from top to bottom, until a rule is found whose e part matches event e' , and whose condition c is satisfied. The action $[o]$ of this successful rule is then defined as the ruling of the law. If no rule succeeds for a given event and state, then the ruling of the law for this event is defined as empty, causing this event to be simply ignored.

Here are some additional details about this informal law-language: (1) the control-state of the home agent can be checked via a predicate of the form $\exists \mathbf{t}$, which means “there exists a term \mathbf{t} in the control state of the home agent”; (2) capitalized symbols represent variables that can match any string; “_” representing a don’t-care symbol; and (3) the phrase $+\mathbf{t}$ is an operation that adds the term \mathbf{t} to the state, while $-\mathbf{t}$ removes such term from the state. These details, along with the comments provided with each example law in this paper should be sufficient for the broad understanding of these laws. Finally, for simplicity, all the examples in this paper deal with groups of agents that operate under a common law, so we do not show any syntax for allowing interoperation between different laws.

Law BC: This law, displayed in Figure 2, consists of three rules, whose effect is described below. Rule $\mathcal{R}1$ of this law is triggered by the `birth` event—the first event in the lifetime of every agent—and its ruling is to add the terms `sBudget(1000)` and `rBudget(2000)` to the *CS* of the newly created agent. The semantics of these terms is defined by the other two rules of this law, as we shall see below.

Rule $\mathcal{R}2$ is triggered by the sending of any message, with the following result: If the message sender has in its *CS* the term `sBudget(B)` with a positive `B`—representing its current sending-budget—this budget would be decremented by 1, and the message would be forwarded to its destination. But if `B` is 0 then the sender would get a “message blocked” response, and its message would not be forwarded. This means that the variable `sB` in the term `sBudget(sB)` in the state of an any *BC*-agent agent x at a given moment in time, represents the upper bounds on the number of messages that x is allowed to send.

Finally, Rule $\mathcal{R}3$ provides the analogous treatment of the receipt of messages, which is limited by the term `rBudget(rB)` of a given *BC*-agent x , representing the current budget of x for receiving messages.

Note the importance of the dual mediation under IC, which allows separate controls over the sending of messages and over their receipt. We will make additional observations about this law in due course.

<pre> $\mathcal{R}1.$ UPON birth DO [(sBudget(1000)), (rBudget(2000))] $\mathcal{R}2.$ UPON sent(.,.) IF \existssBudget(B) AND (B > 0) DO [decr(sBudget(B)), forward] ELSE DO [deliver('message blocked')] $\mathcal{R}3.$ UPON arrived(.,.) IF \existsrBudget(B) AND (B > 0) DO [decr(rBudget(B)), forward] ELSE DO [deliver('message blocked')] </pre>

Fig. 2. Law *BC* of Budgetary Control

3.5 Law Enforcement, and the Concept of Law-Based Trust

By “enforcement” of laws over the communication between distributed actors we do not mean that the actors in question are forced to observe any particular law—indeed it is virtually impossible to force a set of disparate actors, dispersed throughout the Internet, to conform to any given constraint, beyond those implicit in the standard Internet protocols—although it is possible to do so over an intranet, as has been demonstrated in [15]. This section describes what we mean by law-enforcement under the IC model. We start with the definition of a concept of trust, which we call *law-based trust*, or L-trust, for short. In Section 3.5.1 we introduce an example that illustrates some implications of such

trust. In Section 3.5.2 we define our concept of law enforcement, and compare it to the conventional view of enforcement of access control policies.

Definition 2 (law-based trust). *Consider an IC-agent y , and an actor c (which may or may not operate as an IC-agent). We say that c has a law-based trust in y if and only if the following two conditions are satisfied when communicating with y (i.e., getting a message from y , or attempting to send a message to it): (1) c recognizes that y is an IC-agent, and (2) c identifies the law under which y operates. Also, we say that there is a mutual L-trust between two IC-agents if and only if the actor of each of them has an L-trust in the other⁸.*

Of course, the concept of trust over the Internet—as in “trust management” or in “trusted computing base”—is notoriously ambiguous, and somewhat shaky—for a detailed comparison of the conventional concept of trust over the Intranet, and law-based trust (also called “regularity based trust”), see [21]. This is true for our concept of L-trust as well. Nevertheless trust is a necessary concept for the Internet, but it must be carefully justified to be meaningful and dependable. In Section 3.6 we will discuss briefly means for implementing L-trust in a reasonably justifiable manner. Here we turn to the implication of this kind of trust, assuming that it is justified.

The Consequences of L-Trust, and the Concept of \mathcal{L} -Community: The basic consequence of L-trust can be stated as follows. Consider an agent x that operates under law \mathcal{L} , sending a message to an agent y operating under law \mathcal{L}' , which may or may not be equal to \mathcal{L} . The existence of mutual L-trust between x and y has the following consequences: (a) x can be confident that the effect of the arrival of this message at y would be in compliance with law \mathcal{L}' under which y operates; and (b) y can be confident that the message it got from x has been sent in compliance with the law \mathcal{L} , under which x operates.

L-trust is particularly significant for a set of agents operating under a common law \mathcal{L} , which we call an \mathcal{L} -community. Such a common law is necessary the members of an \mathcal{L} -community to collaborate harmoniously and to compete safely with each other—depending on the law that governs them all—even with no knowledge of the nature and intention of each other. This is analogous to the manner in which social laws can create harmonious societies. We will illustrate this phenomena via the following example.

3.5.1 Trustworthy Electronic Theater Tickets—an Example: Consider a theater called *globe* that issues electronic tickets of the form `ticket(d)`, where d represents the date of a performance (assuming there is just one performance a day, and that the number of such tickets issued is equal to the number of seats). The theater then transfers these tickets to its clients, which may be merchants or theater goers. And clients can transfer these tickets to each other any number of times. Finally, the theater admits a client only in exchange of a ticket transferred to it. And suppose that all these transfers of tickets are to be done electronically.

⁸ We sometimes use the term law-based trust whether it is mutual or not, expecting the ambiguity to be resolved by the context.

This scheme can work well—that is, the holder of a ticket can depend on getting a seat in this theater in the specified date—if these tickets cannot be forged. That is, if they can be created only by the *globe* theater, and that once created, they can be transferred, but not copied or modified. We now show how these assurances are provided by a law called \mathcal{TU} , for “ticket unforgeability.”⁹ This law, displayed in Figure 3, consists of three rules discuss below.

Rule $\mathcal{R}1$ enable the theater *globe* (note that the symbol **Self** denotes a variable bound to the id of the home agent) to create any number of tickets in its own control state, simply by sending the message `createTicket(d)`, with any argument `d`. Indeed, the sending of such a message by *globe* triggers this rule, whose ruling is to add the term `ticket(d)` to the *CS* of the sender, i.e., of the *globe* theater. (Note that this law has no rule to match the `birth` event, which mean that state initialization is done, and no control operations are carried out when a \mathcal{TU} -agent is being created.)

<pre> $\mathcal{R}1.$ UPON sent(createTicket(D),_) IF (Self=globe) DO [+ticket(D)] $\mathcal{R}2.$ UPON sent(ticket(D),_) IF (\existsticket(D)) DO [-ticket(D), forward] ELSE DO [deliver('illegal message')] $\mathcal{R}3.$ UPON arrived(_, ticket(D)) DO [+ticket(D), deliver] </pre>
--

Fig. 3. Law \mathcal{TU} of Ticket Unforgeability

The other two rules enable any \mathcal{TU} -agent to transfer a ticket it has to any other \mathcal{TU} -agent¹⁰, simply by sending it a message `ticket(d)`. Specifically, Rule $\mathcal{R}2$, which is triggered by the sending of a message `ticket(d)`, has two possible rulings: If the sender has the term `ticket(d)` in its *CS*, then the ruling would be to remove this term from the *CS*, and to forward this message to its address. Otherwise, the ruling for this event would be to deliver the message “`illegal message`” to the sender, and do nothing else.

Finally, Rule $\mathcal{R}3$, which is triggered by the arrival of a message of the form `ticket(d)` at any \mathcal{TU} -agent, has two operations in its ruling: first, to add the term `ticket(d)` to the *CS* of the home agent (i.e., the agent that received this message), and second to deliver the message itself to the actor of the home agent, to alert it about the arrival of this ticket. (For simplicity, we are not providing here the needed support for the case that the agent receiving such a message is the theater. This would require, in particular, a rejection of a ticket if it is for a wrong date.)

⁹ This law is simplified, in particular, in that it deals only with the transfer of tickets, but does not handle the payment that often (but not always) accompanies such a transfer.

¹⁰ We assume that no interoperability is permitted by this law, and by the other example laws in this paper.

Discussion: The presumed L-trust among the member of this community induces a kind of *communal trust* of the following nature: When a theater *globe* transfers a ticket to some TU -agent, it can be confident that whatever happened to this ticket, there can be no more than one copy of it, anywhere among the members of the TU -community. Correspondingly, when a member of this community receives a ticket from another member, it (or she or he) can be confident that it got the only existing copy of this ticket—regardless of the provenance of the ticket—and thus be confident to be admitted to the theater in the specified date.

This communal trust is particularly remarkable, because the membership of the TU -community is quite open and indefinite. Indeed the TU law imposes no restriction on adopting this law, thus joining this community. Therefore, anybody can join the TU -community at anytime, and there is no way to tell who belongs to it at any given moment in time. (However, as shown in [24], it is quite possible to write laws that can regulate community membership in various ways.) We will consider other aspects of this law below, and in Section 4.

3.5.2 Virtual Enforcement of IC Laws via L-Trust: As has already been pointed out, the IC model does not coerce any actor to exchange \mathcal{L} -messages, under any specific law \mathcal{L} , or to engage in IC-regulated interaction in any other way. Such an engagement is purely voluntary.

Nevertheless, due to L-trust, an actor may often be *virtually compelled* to operate under a particular law \mathcal{L} . Broadly, this is the case when one wishes to use services provided only under this law. For example—returning to law TU above—consider an actor c who wishes to get a ticket for the *globe* theater. This means that one has to get a ticket that is acceptable by the theater, which means, in turn, that the ticket may be obtained either from the theater itself, that operates under law TU , or from some other TU -agent. In either case c would have to operate under law TU to get such a ticket.

Due to this compelling effect of L-trust, we occasionally refer to a mechanism that supports L-trust as a “law enforcement” mechanism. Such, somewhat loose, usage of the term “enforcement” is quite common in the access-control literature. For example, the XACML mechanism [14] relies for its “enforcement” on each server within its purview to voluntarily operate via a *policy enforcement point* (PEP). This PEP is expected to consult a *policy decision point* (PDP) regarding every request sent to it, and then to carry out the decision of the PDP. But no means are provided for forcing servers to operate via such PEPs, and no such enforcement is practical in an heterogeneous system. In fact, the level of enforcement under an IC mechanism is stronger than in conventional access control, due to the existence of L-trust.

3.6 Establishing Law-Based Trust

We start by considering the conditions for a pair of IC-agents to have mutual L-trust in each other, according to Definition 2 of such trust. This is followed by a discussion of how these conditions can be satisfied for large numbers of IC-agents.

Consider a pair $x = \langle c_x, T_x^{\mathcal{L}} \rangle$ and $y = \langle c_y, T_y^{\mathcal{L}'} \rangle$ of IC agents (see Figure 1), exchanging messages under laws \mathcal{L} and \mathcal{L}' respectively—where these laws may be different or equal. The definition of Mutual L-trust between such agents requires the actors c_x and c_y to have L-trust in their counterpart agents. We maintain that such trust would exist if the two actors can trust the pair of controllers $T_x^{\mathcal{L}}$ and $T_y^{\mathcal{L}'}$ to satisfy the following three conditions: (a) each of the controllers operates in compliance with its law, according to the definition of compliance in Section 3.2; (b) the controllers identify to each other the laws under which they operate, via the *one way hash* of these laws (this is sufficient identification if the text of these laws is made public in some well known repository); and (c) the messages exchanged between the controllers, and between them and the actors they serve, are transmitted securely.

A pair of controllers can thus be trusted by the actors that employ them, if they are maintained by a trusted third party (TTP) acceptable to both actors, and if this third party certifies that they satisfy the above three requirements. This assertion is based on a well established tradition in computer research and practice.

Note that it is reasonable to require that the controllers operate on different hosts than the actors served by them, lest they can be corrupted by the actors. However, it may be possible for controllers to reside on the same host with the actors they serve, if the controllers are implemented over TPM [5].

Moreover, to support large numbers of IC agents, which may operate under a variety of laws, one needs to provide a large set of *generic* controllers that can be widely trusted to operate in compliance with any valid law loaded into them. Such a collection of controllers would serve the role of a *trusted computing base* (TCB) for IC-based communication, as suggested by Figure 4. But unlike the traditional TCB, which is usually centralized, our collection of controllers is designed to be decentralized, and is thus referred to as “decentralized TCB”, or DTCB. Such a DTCB needs to be provided by a reliable service, called a *controller service* (CoS), which creates, maintains, continuously tests, and certifies a distributed collection of controller.

Such a CoS can be operated by an enterprise for its own internal use; by a federation of enterprises for the use of its members; or by a regional authority, such as a municipality, for the use of actors in its range. The current implementation of LGI provides an experimental version of a CoS. And there is ongoing research on techniques for protecting the controllers maintained by a CoS from various kinds of attacks; in particular, by intrusion detection, by making it harder to target specific actors, or specific law; and by other means.

It should be pointed out that if an IC mechanism is to be used by arbitrary actors all over the Internet, then the CoS needs to be managed by a reputable commercial company or governmental institution, whose business it is to provide its customers with trustworthy controllers. This organization must be willing to vouch for the trustworthiness of its controllers, and to assume liability for their failures. Such an Internet wide CoS is yet to be established.




Fig. 4. *IC-Based Communication:* Actors are depicted by circles, interacting across the Internet (lightly shaded cloud) via their controllers (boxes) operating under law L . Agents are depicted by dashed ovals that enclose (actor, controller) pairs.

Selective Decentralization of Mediation: Although the mediation of inter-agent interaction under IC is decentralized in principle—each actor operating via its own “private” *controller*—it lends itself, in practice, to a whole range of partially centralized implementations, which can be adapted to different applications.

This is due to the fact that private controllers can be hosted, by the so called *controller pool*, each of which capable of hosting a large number¹¹ of controllers, thus serving several different agents, possibly subject to different laws. The above mentioned CoS is assumed to be a distributed collection of such controller pools.

The availability of controller pools allows one to run any given sub-group of controllers on a single *controller-pool*, thus creating a whole spectrum of possible implementations of mediation—anything between complete decentralization, where each controller of a given community operates on a separate controller-pool, to a complete centralization, where all controllers are hosted by a single controller pool. The potential for performance optimization inherent in such selective centralization of mediation is yet to be fully explored; for some initial ideas about this matter see [24].

¹¹ Under the current implementation of LGI, for example, each controller pool can host several thousands private controllers, depending on the complexity of the law, and on the nature of the application at hand.

4 Analysis of the IC model

The main purpose of this analysis is to identify the main characteristics of the IC model, and to explore the manner in which it satisfies the IC principles we proposed. We will do this, in part, by comparing this model with the traditional access control, attempting to highlight the main structural differences between the two. This analysis will also reveal the underlying reasons for the fact that the IC principles are largely unsatisfied by AC mechanisms.

We start our analysis in Section 4.1, by describing the manner in which scalable sensitivity to the history of interaction is achieved under IC. In Section 4.2 we discuss the global sway that a law can have over the community governed by it—this despite the inherently local nature of laws; and, more generally we argue that the locality of laws under IC does not reduce their expressive power. In Section 4.3 we show that the IC model unifies the DAC and MAC dichotomy of traditional access control; and in Section 4.4 we show that this unification is instrumental in satisfying the principle of separation of policy from mechanism. And we conclude in Section 4.5 by showing how decentralization facilitates flexible, composition-free, interoperation.

4.1 Scalable Sensitivity to the History of Interaction

This section discusses the the manner in which the IC model manages to be sensitive to the history of interaction, as required by our Principle 1, and the nature of the scalability of this sensitivity. We conclude by considering the treatment of these issues under conventional AC.

Generally speaking, history sensitivity is accomplished under IC not by maintaining the complete history of interaction, or even selected parts of this history, in order to be examined when needed. But by maintaining, in the state of each agent x , a relevant *function of the local history* of the interaction of x with other agents. (Note that it is possible, under the IC model, to maintain the entire history of interaction, if necessary, but it is rarely if ever necessary, and it would be very inefficient and unscalable, as we shall see below.)

A simple example of this technique is provided by the budgetary control law \mathcal{BC} introduced in Section 3.4. The purpose of this law is to impose bounds on the number of messages that each agent can send and receive. This is done by (a) initializing the state of each agent with budgets for sending and receiving messages, (b) maintaining the balances of these budgets while each agent exchanges messages, and (c) blocking messages when the corresponding balances are zero. These dynamically changing budgets are the functions of the history of interaction maintained by the law in the state of every \mathcal{BC} -agent.

What makes this technique possible under IC are mainly the following two properties of this model. The first is the fact that the ruling of the law is not restricted to accepting messages or blocking them, but that it can mandate arbitrary change of the local state of the home agent, thus enabling the computation of any desired function of the history of interaction—like the balances of the messaging budgets under law \mathcal{BC} . The second property of IC, which is necessary

for this technique, is the fact that the law can be made sensitive to the state of each agent. The lack of one, or both, of these properties in most conventional AC mechanism, precludes them from providing reasonably scalable sensitivity to history, as we shall see in Section 4.1.2.

4.1.1 On the Scalability of History Sensitivity Under IC: We consider here two quite distinct senses of the term “scalability”: (1) scalability with respect to the length of the process of interaction, and thus length of the history; and (2) scalability with respect to the number of interacting agents, and the volume of messages exchanged between them.

Law \mathcal{BC} is scalable in both of these senses. It is scalable with respect to sense (1) because the incremental computation required for maintaining the balances of the two budgets is small and essentially fixed throughout the interaction history. And it is scalable with respect to sense (2) because the budgetary constraint is defined over the *local* history at each agent. Therefore, this constraint over the behavior of a given agent x depends only on the messaging activities of x itself.

Of course, not every history-sensitive constraint can be implemented under IC in such a scalable manner. One can imagine a non-local constraint, such as the requirements that the total number of messages sent by all members of the community in question does not exceed a certain limit. As argued in Section 4.2, all such constraints can be handled under IC. But the scalability of enforcing such constraints may vary. It is worth noting, however, that many kinds of important history-sensitive constraints are implementable in a very scalable manner under IC. This includes, in particular, the well known *Chinese wall policy* and *dynamic separation of duties* policies, as has been shown in [23] and in [3], respectively.

4.1.2 Sensitivity to History Under Access Control: In a very limited sense, history sensitivity has been part of AC from its inception. Because the distribution of rights among the various subjects, which can be changed by *delegation*, can be viewed as a state of a policy. Moreover, several AC researchers noticed early on that there is a need for additional, but very specific, forms of sensitivity to history, including the so called “Chinese wall” policies and *dynamic separation of duties* policies [29]. But, the first comprehensive attempt at such sensitivity under AC seems to be due to Jajodia et al., in their 2001 paper [16]—a decade after such sensitivity was introduced, essentially under IC, in [22]. But this approach does not scale as it maintains essentially the entire history of interaction in its global set of rules.

The SPL mechanism introduced by Ribeiro et al. [28] is more sophisticated in this respect. It provides an automatic “compilation algorithm” that reduces the amount of historical information to be saved, and attempts to optimize the way that information is to be queried for a given constraint to be evaluated over the history. The technique is clearly more scalable than that of the model of Jajodia et al., but it is not scalable enough. In particular, in order to implement under SPL our simple example of messaging budget, with initial budget of N , one will have to keep at least part of the history of up to N messages—where N is

the initial messaging budget—and this history will have to be kept in a central place, which makes it quite unscalable for large values of N .

The underlying reasons for this lack of scalability of SPL, and of Jajodia’s model, are that both of them retain the traditional AC model, with its centralized enforcement mechanism; and under which a policy can only rule to permit or deny the transfer of a message, but cannot also change any state, as is done under law \mathcal{BC} above.

4.2 On the Global Sway, and Expressive Power, of IC Laws

The local nature of IC laws raises the following questions: (a) can such laws establish global properties of a system consisting of many agents; and (b) is this ability to establish global properties as general as that of laws that are not restricted to be local? The answer to both of these questions turns out to be affirmative. We will address these issues here for an \mathcal{L} -community, that is, for a set of agents that operate under a common law \mathcal{L} . But our conclusions are valid more generally. Particularly under hierarchical organization of laws, described in [2].

We start with some terminology. We will refer to a global property of a multi-agent system as *regularity*, by which we mean “conformance [by every part of the system] to a given rule.” And we distinguish between two kinds of regularities of a given multi-agent system S , which we call *local* and *non-local*. A local regularity of system S can be expressed as a universally quantified formula of the form $\forall x \in S P(x)$, where $P(x)$ is a local property of agent x ; that is, $P(x)$ is defined over the local state of x , and over local events at it. For example, $P(x)$ can be a given limit of the number of messages that an agent x can send (as established by law \mathcal{BC} of Section 3.4).

A non-local regularity is one that cannot be expressed in this manner. For example, consider the following statement: “*the aggregate of all messages sent by all members of system S does not exceed 1000.*” This is a regularity in that it restricts the behavior of each member of S —in particular, no agent is allowed to send the 1001th communal message. And this regularity is non-local in that this constraint on message sending depends on the coincidental state and behavior of all members of S .

Now, local regularities are the direct consequences of the IC model, due to the fact that all members of an \mathcal{L} -community conform to the same law \mathcal{L} . This has been demonstrated by our example laws \mathcal{BC} in Section 3.4, and by law \mathcal{TU} in Section 3.5.1.

The case of non-local regularities is less obvious. Clearly, non-local regularities can be established by a non-local law enforced via a central reference monitor that mediates all messaging activities of members of S , and can maintain the global control state of the entire system (like the number of messages sent so far by all members of S .) The question is: can non-local regularities be established via the local IC laws?

The affirmative answer to this question has been provided in [25], via the following result:

Any non-local constraint on interaction that can be implemented via a central reference monitor can be localized, and thus implemented via a local IC law.

This important result—whose proof is quite straightforward, but too long to be repeated here—is due to the ability of IC laws to mandate the movement of relevant information from one agent to another, getting such information to the appropriate decision points. (Unfortunately, space limitation does not permit even an example of how this is done, but see [33] for such an example).

This ability to localize non-local regularities means, essentially, that *the locality of IC laws does not reduce their expressive power*. Of course, such localization involves certain amount of aggregation of information, which may reduce its scalability. But we have found that in practice, many useful non-local constraints involve modest amounts of aggregation, which has only a marginal effect of scalability. For more about this matter, see [24].

In conclusion, we can state that the locality of laws does not impede their ability to establish global properties, and it does not reduce their expressive power.

4.3 The Unification of Mandatory and Discretionary Controls

We discuss here a critical structural aspect of the IC model: its unification of mandatory and discretionary aspects of control, which are generally considered independent elements of access control. Indeed, the 1999 report “Trust in Cyberspace,” by the National Academy of Science, characterizes the conventional access control technology as follows ([30] page 114): *the two basic types of access control policies that have dominated the computer security work for over two and a half decades [are the] discretionary access control [DAC] and mandatory access control [MAC]*. DAC policies, this report continues to explain, are discretionary in the sense that a subject can pass (via delegation) its access permissions on to another subject; and MAC policies enforce access control on the basis of fixed set of rules mandated by a central authority. (Note that MAC is frequently identified with *multi level security* (MLS) policies, like that of Bell-Lapadula [8], but we adopt here the general interpretation of MAC). In other words, DAC and MAC are viewed as a dichotomy, i.e., as two independent, and distinct types of policies. And an AC mechanism may support one, or the other, or both—in particular RBAC is claimed [27] to support both, (although that paper takes the narrow interpretation of MAC, as an MLS policy).

The Problem with the MAC/DAC dichotomy: We maintain that the MAC/DAC dichotomy is problematic from both the theoretical and practical points of views. The theoretical problem with this dichotomy is that every DAC policy gets its semantics from an underlying MAC policy. Consider for instance a *capability-based AC*—a common example of a “pure DAC policy”. The very meaning of capabilities, as well as their origin, delegation and revocation, are all defined by an essentially mandatory (MAC) policy built into the very fabric

of the AC mechanism at hand, and are not subject to the discretion of the subjects. It is this *implicit nature* of the mandatory counterpart of the discretionary distribution of capabilities that makes it seem to be a pure DAC, which it is not.

The practical consequence of this dichotomy is that it encourages making the mandatory counterpart of what is viewed as a pure DAC policy implicit in the mechanism, as we have seen above. This reduces the flexibility and generality of AC mechanisms, because any change in this implicit mandatory policy—like a change of the delegation process—would involve changing the mechanism itself, or replacing it with another. We will return to this matter in the following section.

The Unification of MAC and DAC Under the IC Model: The IC model is based on two complementary and inextricably intertwined elements: (a) the explicitly stated fixed *law*, which is mandatory for all members of the community in question; and (b) the *state* that can be changed at the discretion of the various actors, subject to the law. So, while the law and the state correspond roughly to the concepts of MAC and DAC, they cannot be viewed as a dichotomy, but are, rather, the elements of a unified control mechanism. Indeed, the ruling of the law depends on the state; and the state changes in response to events that occur mostly at the discretion of some actors, but only according to the mandatory rules laid out by the law. They are, thus, unavoidably intertwined. To illustrate these points we introduce below an implementation under IC of a version of capability based access control.

4.3.1 Capability Based Access Control Under the IC Model—an Example: We introduce here a law, called \mathcal{CB} (for “capability based”), which establishes the following, informally stated, provisions. First, under this law a capability is a term $cap(y, d)$, which, when contained in the CS (i.e., control state) of some agent x , signifies the right of x (its holder) to send messages of the form $msg(m)$ (with any content m) to agent y . Second, a capability $cap(y, d)$ is *delegatable* only if d is 1. And the delegation of a $cap(y, 1)$ capability to any agent z can be done by the holder of this capability, by sending to z the message $delegate(cap(y, d'))$, where d' can be either 1 or 0. The arrival of this message at z would provide z with a delegatable or an undelegatable version of this capability, depending on whether d' is 1 or 0. (Note that this is just one of many possible ways for doing delegation; we will return to the general topic of delegation later.) Finally, every new member x of the \mathcal{CB} -community is initialized with the delegatable capability $cap(x, 1)$ *for itself*—which according to the above, it would be able to delegate to any other member.

A law \mathcal{CB} that specifies this policy under the IC model is displayed in Figure 5. This law consists of five rules. Rule $\mathcal{R}1$ is triggered by the *birth* event, and its ruling is to add to the state (CS) of the home agent a term $cap(\text{Self}, 1)$ representing a delegatable capability for itself. This means that every agent operating under this law would be initialized in this way.

Rule $\mathcal{R}2$ is triggered by the sending of any message of the form $msg(_)$ to any agent Y . This rule produces the ruling to forward the sent message only if the condition $\exists cap(Y, _)$ is satisfied; that is if the sender has a cap term for Y

$\mathcal{R}1.$	UPON birth DO [+cap(Self,1)]
$\mathcal{R}2.$	UPON sent(msg(_),Y) IF $\exists \text{cap}(Y, _)$ DO [forward] ELSE DO [deliver('illegal message')]
$\mathcal{R}3.$	UPON arrived(_,msg(_)) DO [deliver]
$\mathcal{R}4.$	UPON sent(delegate(cap(Y,D)),_) IF ($\exists \text{cap}(Y,1)$) DO [deliver]
$\mathcal{R}5.$	UPON arrived(_,delegate(cap(Y,D))) DO [+cap(Y,D), deliver]

Fig. 5. Law \mathcal{CB} of Capability Based Control

in its state. Otherwise, this rule would produce a ruling to deliver to the sender an “illegal message” warning. When a message sent under this rule arrives at its destination it is delivered without further ado by Rule $\mathcal{R}3$. Note that these two rules provide the $\text{cap}(\dots)$ terms with the semantics of a *right to send messages*, i.e., the semantics of capabilities.

Rule $\mathcal{R}4$ is triggered by an agent sending a message $\text{delegate}(\text{cap}(Y,D))$ to any agent Z . The ruling for such sending would be to forward this message, but only if the sender has the term $\text{cap}(Y,1)$ in its CS , representing a delegatable capability. Finally, when this delegate message arrives at its destination Z , it would trigger Rule $\mathcal{R}5$ at Z , which would cause the sent capability to be added to the CS of the receiver.

Discussion: This example clearly demonstrates the manner in which mandatory and discretionary aspects of regulation are represented under IC by means of the law and the state, respectively; and the manner in which these two elements of IC are intertwined. It is law \mathcal{CB} that defines the semantics of capabilities, their initial distribution, and the manner in which they can be delegated. But the eventual distribution of the capabilities in the state of the various agents is determined by the delegate messages sent at the discretion of the various actors, but subject to the law.

Moreover, the regime defined by law \mathcal{CB} is only one of many types of capability-based regimes that can be implemented under IC. One can, in particular, replace \mathcal{CB} with a law that employs a different initialization of the states of agents, and a different manner for delegating capabilities. The significance of this flexibility of the IC model is discussed in the following section.

4.4 Separation of Policy from Mechanism

As has already been pointed out, one of the original principles of secure computing has been that of *separation of policy from mechanism*, formulated in the early 1970s, by Wulf et al. [32], and by others. What this principle means, essentially, is that an AC mechanism should be as general as possible, without a bias toward

any particular kind of policies, thus accommodating a wide range of them. But despite the obvious importance of this principle—particularly for complex systems that need to use multiple policies—it has not been satisfied to a significant degree by conventional AC mechanisms. (The importance of this principle has been recently noted in [11], but without proposing a significant support for it.) This, we claim, is mostly because of the failure of these mechanisms to make the mandatory aspects of a policy explicit.

The situation is very different under the IC model because its law represents mandatory constraints explicitly. The lack of bias of this model toward any particular type of policy is evident from our \mathcal{CB} law. Although the IC model has no built-in concept of capability, this law creates a certain interpretation of these concepts, as it can easily create others. We now demonstrate the degree to which IC satisfies the principle of separation of policy from mechanism, via brief case studies of the treatment of two well known aspects of control—*delegation*, and *roles*—under conventional AC mechanisms, and under IC.

4.4.1 Delegation—a Case Study: The basic concept of delegation is very simple: it is the giving to somebody else a right that one has for oneself. But there are various possible ways for delegation to be carried out, and various kinds of constraints that one may want to impose on the process of delegation. For example, one may specify such things as who is permitted to delegate which privileges to whom, and under which circumstances; and how such delegation should affect the privileges of the delegator itself. We call a specific choice among such possibilities a *delegation-pattern*. And here is a sample of possible patterns of this kind: (1) The requirement that certain rights shall be delegated only *by transfer*; that is, the delegator will give up its own right, when delegating it to somebody else, as we have done with theater ticket under law \mathcal{TU} . (2) Some restriction on the number of times a given right may be delegated, if at all. (3) Making delegation of a privilege operate as loans, requiring the privilege to be returned within a specified time period. And (4) Establishing mechanisms for the revocation of delegated privileges.

Despite the need of different applications to employ different delegation patterns, most conventional AC mechanisms employ a fairly simple, and quite rigid, built-in pattern. According to [6], “*The standard approach to delegation is binary: either delegation is possible, and then no substantial further control over the way it is used is possible, or else no delegation is permitted.*” Recently, several more sophisticated delegation patterns have been proposed. In particular, Bandmann et al. [6] proposed a model of *constraint delegation* which allows one to control the possible shapes of delegation chains. Also, Li et al. [19] permitting a fixed upper bound to be imposed on the depth of delegation chains. But each of these delegation patterns requires a new AC mechanism to be constructed—contrary to the principle of separation of policy from mechanism.

Under the IC model, on the other hand, the nature of delegation is specified by the law of a community, as it is the case with law \mathcal{TU} of Section 3.5.1, and law \mathcal{CB} of Section 4.3. And it is easy to see how other types of delegation patterns can be similarly defined. For a demonstration of how sophisticated delegation

that combines several of the above mentioned patterns can be defined by a law under an IC mechanism, see [4, 1].

4.4.2 Role Based Access Control (RBAC)—a Case Study: An important extension of the traditional matrix-based AC model has been the introduction of the “role-based access control” (RBAC) model, by Sandhu and his colleagues [29]. This model views roles—which represent such things as the position that a given individual holds in a given organization—as a set of permissions, and defines an AC-policy essentially as a pair of mappings: a mapping of users to sets of roles, and a mapping of roles to sets of permissions. The RBAC model also provides for a hierarchical organization of roles, and includes certain constraints on the combinations of roles that a given individual can possess at one time.

The RBAC model has been widely adopted by the industry and been broadly accepted by the research community. But despite its considerable structural complexity (the official RBAC has four “reference models” of increasing sophistication) RBAC proved to be overly rigid, which gave rise to a host of *RBAC-based models* that generalized RBAC in various ways. These variants of RBAC include, but are not limited to: (1) parameterization of the roles of a subject by its other attributes [13]; (2) making authorization dependent on some dynamic relationship between the subject, and the object on which it operates [7]; (3) making the rights associated with a role dependent on the *context*, such as time and location [26]; and (4) providing for dynamic, and temporal, control over role-permission activation and deactivation, and over user-role assignment and activation (TRBAC [9] and GTRBAC [17]).

Now, all these models are worthy generalizations of the original RBAC. The problem is that each of them requires a different AC mechanism for its support—again, in a sharp breach of the principle of separation of policy from mechanism.

Under IC, on the other hand, we have shown in [3] that the original RBAC model itself, and every one of the above mentioned extensions of it, can be defined simply by writing a suitable law. And, thus, they can all be realized via a single IC mechanism. And this, despite the fact that IC has no built-in concept of roles.

4.5 Flexible, Composition-Free, Interoperability

We define interoperability under IC as the ability of two agents operating under different laws to interact, without violating their own rules. Although this is analogous to interoperability under access control, the mechanisms of interoperability under these two control regimes is substantially different.

The conventional AC approach to interoperability [20] between parties operating under policies $P1$ and $P2$, has been to *compose* these policies into a single policy $P12$, which is, in some sense, consistent with both $P1$ and $P2$. The composition $P12$ is then to be fed into an appropriate reference monitor, which would mediate the interaction between the two parties. But composition of policies has several serious drawbacks: manual composition is laborious, and error prone; and automatic composition is computationally hard [20], and often

impossible. Yet, composition is the natural, and perhaps necessary, approach to interoperability under AC—because AC employs a single reference monitor to mediate the interaction of any pair of agents.

Under the IC model, on the other hand, composition of laws is neither natural nor necessary. This is because IC employ dual mediation for every pairwise interaction. This enables an $\mathcal{L}1$ -agent $x1$ to exchange messages with an $\mathcal{L}2$ -agent $x2$, if such an exchange is permitted by both laws $\mathcal{L}1$ and $\mathcal{L}2$, and without the need to form a composition of $\mathcal{L}1$ and $\mathcal{L}2$. The basic mechanics for such interoperation is described in [31], where we have shown how it can be utilized for regulating B2B transactions.

Of course, for a law $\mathcal{L}1$ to permit the exchange of some messages with agents operating under $\mathcal{L}2$, it—or, rather, its author—should examine the other law, or otherwise have some trust in it. This is not necessary if the IC mechanism at hand supports the concept of *law hierarchy*, which is one of the possible extensions of the IC model, introduced by the LGI mechanism [2]. In a nutshell, if laws $\mathcal{L}1$ and $\mathcal{L}2$ are both hierarchically *subordinate* to a common law \mathcal{L} , then they are guaranteed to conform to \mathcal{L} . This fact turns out to be sufficient for a very flexible and efficient composition-free regulation of dynamic coalitions of organizations (such as grids and virtual-organizations), as has been shown in [2].

5 Conclusion

We have proposed here five principles that an interaction control mechanism needs to satisfy in order to support an effective governance of heterogeneous distributed systems. And we have introduced a minimalist reference model of control mechanisms that satisfy these principles. This model is an abstraction of the existing *law governed interaction* (LGI) mechanism, and it is intended to serve as a basis for a whole family of concrete *IC mechanisms* that would satisfy the IC principles, but may be designed for different types of communication, different performance requirements, and different application domains; and which may extend the minimalist IC model in various ways. One such mechanism, for ad hoc wireless communication, has been recently defined and implemented [12].

The IC model differs from conventional *access control* (AC)—the currently dominant approach to regulation of distributed systems—in its objectives, its underlying structure, and its resulting properties. The main objective of conventional AC is to permit access only to those that have the right for it, but without any concern about the dynamic behavior of those that do get access. The IC model, on the other hand, aims to regulate the dynamic behavior of interacting parties as well.

Structurally, IC differs from conventional AC, mainly in two ways: First, IC is inherently decentralized, although it lends itself, in practice, to a whole range of partially centralized implementations. Second, IC replaces the concept of AC *policy*—traditionally defined as an *access control matrix*—with the pair $\langle \text{law}, \text{state} \rangle$, where the fixed law holds sway over an entire community, while the state is a distributed collection of local states, each of which can change dynamically independently of the others, but subject to the global law. Among

other things, this device unifies the concepts of DAC and MAC, conventionally viewed as a dichotomy under AC.

The consequences of these structural differences between the IC model and conventional AC are quite profound. Besides being far more general than access control, the IC model is substantially more scalable, lends itself to generalizations, and supports the principle of “separation of policy from mechanism”. These differences between the IC model and AC, and some others, have been studied in some detail in this paper.

Finally, it should be pointed out that the IC-model itself has not been implemented due to its abstract nature. However, the LGI mechanism, which the IC model is an abstraction of, has been implemented and released [24]; and it has been applied experimentally to a wide range of applications, some of which have been cited in this paper.

Acknowledgment: I am indebted to Yaron Minsky for many rewarding discussions on this topic.

References

1. X. Ao, N. Minsky, and V. Ungureanu. Formal treatment of certificate revocation under communal access control. In *Proc. of the 2001 IEEE Symposium on Security and Privacy, Oakland California*, May 2001.
2. X. Ao and N. H. Minsky. Flexible regulation of distributed coalitions. In *LNCS 2808: Proc. European Symp. on Research in Computer Security (ESORICS)*, Oct. 2003.
3. X. Ao and Naftaly H. Minsky. On the role of roles: from role-based to role-sensitive access control. In *Proc. of the 9th ACM Symposium on Access Control Models and Technologies, Yorktown Heights, NY, USA*, June 2004.
4. Xuhui Ao and Naftaly H. Minsky. Regulated delegation in distributed systems. In *Proc. of the IEEE 7th International Workshop on Policies for Distributed Systems and Networks, London, Ontario, Canada*, June 2006.
5. W Arbaugh, D Farber, and J Smith. A secure and reliable bootstrap architecture. In *Proceedings of 1997 IEEE Symposium on Security and Privacy*, 1997.
6. O. Bandmann, M. Dam, and S. F. Babak. Constrained delegations. In *proceedings of 2002 IEEE Symposium on Security and Privacy*, May 2002.
7. J. Barkley, K. Beznosov, and J. Uppal. Supporting relationships in access control using role based access control. In *Proceedings of the Fourth ACM Workshop on Role-Based Access Control*, pages 55–65, October 1999.
8. D.E. Bell and L.J. LaPadula. Secure computer systems: Mathematical foundations and model. Technical Report M74-244, Mitre Corp., August 1975.
9. E. Bertino, P.A. Bonatti, and E. Ferrari. Trbac: A temporal role-based access control model. *ACM Tran. on Information and System Security*, 4(3), 2001.
10. M. Blaze, J. Feigenbaum, J. Ioannidis, and A. Keromytis. The keynote trust-management systems, version 2. ietf rfc 2704. Sep 1999.
11. M. Coetzee and J.H.P. Eloff. Virtual enterprise access control requirements. In *Proceedings of SAICSIT 2003 Conference*, pages 285–294. ACM, October 2003.

12. Rishabh Dudheria, Wade Trappe, and Naftaly Minsky. Coordination and control in mobile ubiquitous computing applications using law governed interaction. In *Proc. of the Fourth International Conference on Mobile Ubiquitous Computing, Systems, Services and Technologies (UBICOMM) Florence, Italy*, October 2010.
13. L. Giuri and P. Iglio. Role templates for content-based access control. In *Proc. of the ACM Workshop on Role-Based Access Control (RBAC'97)*, 1997.
14. S. Godic and T. Moses. Oasis extensible access control. markup language (xacml), version 2. Technical report, Oasis, March 2005.
15. Z. He, T. Phan, and T.D. Nguyen. Enforcing enterprise-wide policies over standard client-server interactions. In *Proc. of the Symp. on Reliable Distributed Systems (SRDS)*, 2005.
16. S. Jajodia, P. Samarati, M. L. Sapino, and V. S. Subramanian. Flexible support for multiple access control policies. *ACM Trans. on Database Systems*, 26(2), June 2001.
17. James B. D. Joshi, Elisa Bertino, Basit Sahfiq, and Arif Ghafoor. Dependencies and separation of duty constraints in grbac. In *Proceedings of the 8th ACM Symposium on Access Control Models and Technologies (SACMAT 2003)*, 2003.
18. B.W. Lampson. Protection. In *Proceedings of 5th Princeton Symposium on Information Sciences and Systems*, pages 437–443, March 1971. Reprinted in *ACM Operating Systems Review*, Vol 8(1), pp 18-24 (Jan. 1974).
19. Ninghui Li, Benjamin N. Grosf, and Joan Feigenbaum. Delegation logic: A logic-based approach to distributed authorization. *ACM Transaction on Information and System Security (TISSEC)*, pages 128–171, Feb 2003.
20. P. McDaniel and A. Prakash. Methods and limitations of security policy reconciliation. In *Proc. of the IEEE Symp on Security and Privacy*, May 2002.
21. N. H. Minsky. Regularity-based trust in cyberspace. In *In LNCS 2692: the Proceedings of the First International Conference on Trust Management, Crete, Greece*, May 2003.
22. Naftaly H. Minsky. The imposition of protocols over open distributed systems. *IEEE Transactions on Software Engineering*, February 1991.
23. Naftaly H. Minsky. A decentralized treatment of a highly distributed chinese-wall policy. In *Proc. of the IEEE 5th Int. Workshop on Policies for Distributed Systems and Networks, Yorktown Heights, NY, USA*, June 2004.
24. Naftaly H. Minsky. *Law Governed Interaction (LGI): A Distributed Coordination and Control Mechanism (An Introduction, and a Reference Manual)*, February 2006. (available at <http://www.moses.rutgers.edu/>).
25. Naftaly H. Minsky and V. Ungureanu. Law-governed interaction: a coordination and control mechanism for heterogeneous distributed systems. *TOSEM, ACM Transactions on Software Engineering and Methodology*, 9(3):273–305, July 2000.
26. M.J. Moyer and M. Abamad. Generalized role-based access control. In *Proc. of the 21st Intern. Conf. on Distributed Computing Systems*, pages 391–398, 2001.
27. S. Osborn, R. Sandhu, and Q. Munawer. Configuring role-based access control to enforce mandatory and discretionary access control policies. *ACM Transactions on Information and System Security*, 3(2):85–106, May 2000.
28. C. Ribeiro and P. Ferreira. A policy-oriented language for expressing security specifications. *International Journal of Network Security*, 5(3), November 2007.
29. R.S. Sandhu, D. Ferraiolo, and R. Kuhn. The nist model for role-based access control: Towards a unified standard. In *Proceedings of ACM Workshop on Role-Based Access Control*. ACM, July 2000.
30. F.B. Schneider, editor. *Trust in Cyberspace*. National Academy Press, 1999.

31. V. Ungureanu and Naftaly H. Minsky. Establishing business rules for inter-enterprise electronic commerce. In *Proc. of the 14th International Symposium on Distributed Computing (DISC 2000); Toledo, Spain; LNCS 1914*, October 2000.
32. W. Wulf, E. Cohen, W. Corwin, A. Jones, C. Levin, C. Pierson, and F. Pollack. Hydra: The kernel of a multiprocessor operating system. *CACM*, 17:337–345, 1974.
33. Wenxuan Zhang, Constantin Serban, and Naftaly H. Minsky. Establishing global properties of multi-agent systems via local laws. In Danny Weyns, editor, *Environments for Multiagent Systems III, LNAI 4389*. Springer-Verlag, 2007.