To appear in the 1996 *IEEE Symposium on Research in Security and Privacy,* Oakland. [Pre-final Version]

A Communication Agreement Framework of Access/Action Control

Martin Röscheisen and Terry Winograd
http://pcd.stanford.edu

Computer Science Department
Stanford University

Abstract

We introduce a framework of access/action control which shifts the emphasis from the participants to their relationship. The framework is based on a communication model in which participants negotiate the mutually agreed-upon boundary conditions of their relationship, and create social reference points by encapsulating them in compact "communication pacts," called "commpacts." Commpacts are designed to provide a language enabling a social mechanism of coordinated expectation. We argue that in networked environments characterized by multiple authorities and "trusted proxies," this model can deal with the complexities of general (user- and content-dependent) distributed access/action control and provides a clear user-conceptual metaphor. The framework embeds naturally into the existing legal and institutional infrastructure; it generalizes work in electronic contracting. Commpacts can be seen as a third fundamental type next to access-control lists (ACLs) and capabilities.

Keywords: Access Control, User-Conceptual Models, Networked Environments, Trusted Proxies, Rights Management Systems, Privacy, Electronic Contracting.

1 Introduction

Designs are significantly affected by the assumptions on their intended context of use. This is why access control design for applications with "military security" requirements lead to quite different security models (e.g. Biba [1]) than designs targeted at a certain other, predominantly "commercial" setting—such as the model which Clark & Wilson [2] arrive at by emphasizing integrity over secrecy. Clearly, such requester-granter focused "protection" has yet different underlying assumptions than those in the case of, say, someone making an electronic newsletter available on the Internet to student subscribers, a library providing a site license, or an online community keeping up a shared space.

In this paper, we present a model of access/action control which is based on interaction partners coordinating their expectations by articulating and negotiating the mutually agreed-upon boundary conditions of the kinds of relationship they are willing to maintain; access control occurs then as ancillary to such relationship management.

We first look briefly at the structural characteristics of the kinds of networked environments we are interested in. In particular, we identify the assumption of having certain organizations ("home providers") serve as "trusted proxies." We then place access control in context of an abstract system architecture, and review some basic properties of generic models. In particular, we identify high negotiation complexity as a constraint on usability and feasibility of general (user- and content-dependent) distributed access control, and we look in detail at one such example.

We then introduce a model of access control that is framed in a subject-subject communication model where the boundary conditions of the relationships are tokenized in a first-class-citizen object. This object represents a (relational) "contract" between interacting persons (a "communication pact," "commpact"). It is an encapsulation that enables social reference point creation as part of a language for coordinating expectations. We argue that the model is thus suited to handle action control in networked environments which exist within a much larger social, economic, and legal framework than its predecessors were designed for.

Specifically, we argue that the commpact model provides a natural metaphor for a uniform user-conceptual model of action ("access" as well as "usage") control in a complex environment. It conceptually matches with the existing legal infrastructure, which builds upon a subject-subject world with contractual relations, and fits into the institutional infrastructure to which we are exposed as the context of the technical design. In particular, a trusted-proxy networked environment helps in finding a useful implementation model.

Commpacts provide for usability by encapsulating interdependent authorization policies and factoring out unintended constraint interactions; they supersede the problem of generally high negotiation complexity by localizing authorization interactions which are interrelated according to some usage context.

This paper focuses on laying out the basic framework; details of a prototype system which is being built based on this framework as part of the Stanford Integrated Digital Libraries project can be found in [24]; this includes specifications of the commpact language, the object-request protocols, enforcement issues, and a range of example scenarios.

2 Networked Environments with Trusted Proxies

Let us briefly describe how we view the overall structure of the kinds of emerging and future networked environments towards which we are interested in targeting our work for. In particular, as argued in [31][32], this includes the presence of what we call "home providers," which serve as trusted proxies next to clients and servers (cf. Figure 1).

Home Providers

The current "online services" like AOL and Compuserve are examples of preliminary versions of such home providers for the case of consumer presence on the Internet. Universities and companies currently provide similar in-house services for students and employees; each of these can be seen functionally as just another instance of the current forms of a home provider. In particular, Winograd [32] argues that in the near future the bulk of accesses on the Internet will take place via such "proxies."

But the notion of a home provider also extends to other domains. In the case of electronic trading (using EDI standards; cf. [9] for an introduction), what is called here an "EDI network provider" fulfils quite clearly functions which we attribute more generally to the notion of a home provider (e.g. reputation-based management of membership, authenticating members, certifying user attributes, etc.). In the case of the local-loop phone networks, current phone companies like Pacific Bell can be seen as being home providers for interactions conducted on these networks. And in the case of credit processing networks, companies like VISA act as home providers for certain financial transactions.[1]

Generally, home providers have service contracts with their members (e.g. the online service provider contract, the "EDI trading partner agreement," etc.) that allow them to regulate which kinds of electronic activities are binding under which terms and conditions, and to which rules their members are committed for interactions among themselves. For example, these agreements might stipulate that digital signatures are considered legally binding, that logged data counts as evidence, or that whenever someone replies in a certain way to a certain message type, then a contract has been formed [35][34]. Economically speaking, these organizations are generally instances of reputation-based community enforcement institutions (cf. generally Milgrom and Roberts [19]); they realize transaction cost efficiencies by integration.

1. While there is speculation that proprietary (value-added) networks of the kinds mentioned will go up into a more general internet, this primarily affects the lower-level technical protocols, not the institutional structures (which are likely to remain).

Trusted Proxies

Architecturally, home providers can be characterized as trusted third-party organizations which operate as *proxies*, that is, they are a third party with the special property that all requests go through this party one way or the other (and as a corollary, we can say that there is *no overhead* in including such a trusted third party into interactions between clients and servers).

FIGURE 1. Networked Environments with Trusted Proxies.

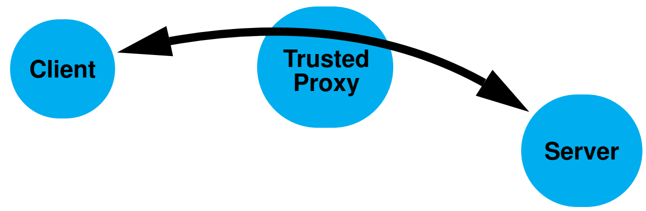

We view the notion of a trusted proxy as a structural characteristic of networked environments which we are interested in considering in the design of new technological solutions. Indeed, as we will see later, the communication agreement model we propose leverages from and fits naturally into such an environment—more readily than would be obvious under different assumptions on the environment (where it might have appeared uninteresting).

3 Access Control in Context

Let us look at how access control fits into an overall architectural picture. The ISO Access Control Framework [6] places access control in the context of an abstract technical system, and introduces relevant terminology (cf. Figure 2): Every operation is intercepted by an "Access Control Enforcement Facility" (AEF), which checks with an "Access Control Decision Facility" (ADF) whether this operation is admissible. If so, the action is performed on the target object; otherwise a failure exception occurs.

The access-control decision facility itself is now generally conceived to be based on a system of access-control rules in which specific policies are expressed (e.g. LaPadula [5]). These rules will in turn generally depend on properties of the system context (e.g. the time), the target object (content-dependent access control), and the requester (user-dependent access control). An example for such a general policy which we might want to accommodate could be "Approve all requests from US citizens for documents which have not been modified since last week."

Note that in a networked environment the ADF itself will generally be distributed, and requests from the target's trusted reference monitor for confirmation of user attributes (cf. "Check with User" in Figure 3) would itself be intercepted by access control again—this time by the user's trusted reference monitor. Such trusts will usually not extend to the same reference monitors. Also note that in a

FIGURE 2. ISO Access Control Framework.

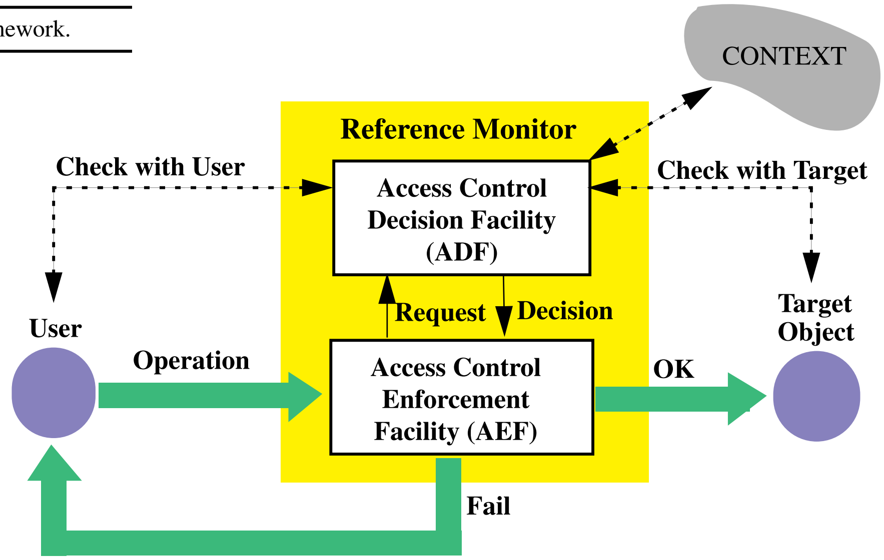

distributed setting each participant has their own authority to determine by which rules they wish to participate in the system.

this general case, the reference monitor takes on then more the role of a "negotiator" between client and target. Although simplifications are possible in the case where certain entities are fully trusted, in general the rule interactions within a given reference monitor and those between different reference monitors (user-trusted, target-trusted) are less than obvious, and the negotiation complexity can easily get intricate.

Not only must the access rules within one policy module be appropriate, but they also have to work together in the right "incremental revelation" schedule with those of the rules of other modules. For example, in the nationality-based policy mentioned above, the requester has to understand that when claiming access with respect to this policy, then the otherwise private nationality attribute must be revealed to the target's reference monitor.

Such interactions will often become quite complex, difficult to devise, debug, and understand; thus Moffet and Sloman [18] conclude that such general, application-independent access control will therefore not be practical.

Example: Negotiation Complexity for General CallerID Interactions

Let us consider here as a simple demonstration example a set of rules by which people might want to control access to their attention/presence for phone interactions, that is, access control to the phone bell. This is an example for which the privacy implications of a specific choice of such access rules have been extensively debated under the name "CallerID"; it is also an example where we believe much can be resolved by going beyond the limited nature of the access control provided in the considered communication systems—to a general access-control system which enables participants to articulate the conditions under which they are willing to participate in a communication exchange.

FIGURE 3. Interactions between Partially Trusted, Distributed Authorization Modules.

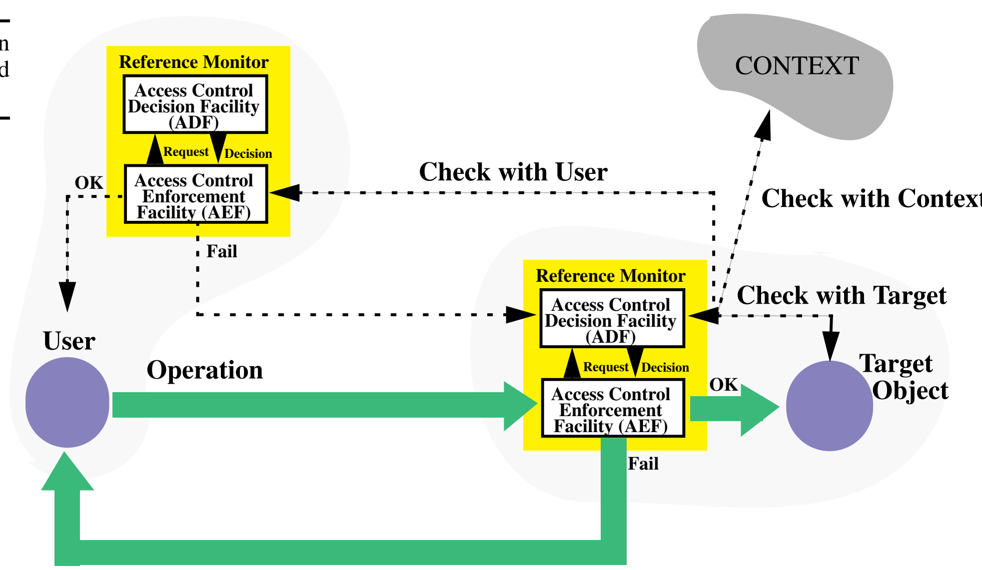

FIGURE 4.
CallerID Example: Simple Set of Phone-Access Rules.

Person A: Tom

```
name='Tom'.
ID='72355'.

reveal(name):-
 B.callType=
 'private'.

reveal(ID).
```

Person B: Lisa

```
name='Lisa'.
callType='private'.

reveal(name):-isFriend(A).
reveal(callType).

connect_call:- NOT
 block_call AND good_call.

block_call :-
  Context.time='evening'
  AND
  (A.ID='SJMN' OR NOT A.ID)

good_call :- isFriend(A).

isFriend(A):-A.name='Tom'.
```

Consider two communication participants Tom and Lisa, each of which expresses their preferences in a set of access-control rules (cf. the pseudo datalog [26] in Figure 4).

Each person has a set of attributes such as `name`, `ID`, and `callType`, which are communicated only when the corresponding `reveal` access predicate allows it. The phone bell is accessed here by the function `connect_call`, which determines whether or not the bell is rung.

A notation of `A.name` is used to access the `name` attribute of A; if this is called by someone other than A, then A is asked to reveal this attribute. Specifically, `reveal` is here a special predicate about a personal information attribute; if there is a rule which makes it true, then the corresponding attribute is returned.[2] Note that privacy considerations dictate that only the least amount of information needed for accomplishing a certain task should be released under any given circumstances.

Figure 5 lays out the temporal sequence (top-down) of how the authorization policies interact when A calls B in the evening. Notice the brittleness of the system: With the rule authorities distributed, bugs can easily be introduced by one side by not sufficiently considering the possible dynamics which might result from unexpected interactions with unknown policies at other sites. Indeed, in the general case, not even the possibility of deadlocks can be ruled out. Not only is the negotiation cost high here in terms of network transactions, but, more significantly for our analysis here, the interaction interdependencies are quite intricate and unclear, and the usability of such a system is therefore likely to be low.

The underlying reason for this is of course that what we have here is a coordination problem which needs a language shared among the participants; if such a language provides high-level primitives to coordinate and to provide context for any necessary lower-level transactions, then much of the negotiation complexity can be simplified. Note that A does not know a priori what B wants to know, and vice versa.

The conceptualization which we suggest as an access control framework is targeted at avoiding this interdependency/negotiation complexity by bridging the gap between requester and target with an intermediate concept which encapsulates access control policies that "belong together." This would reduce negotiation complexity and might enable the kinds of general access control policies which we would like to have.

4 Understanding Conventional Access Control Models

Let us revisit the conventional textbook description of the fundamentals of access control: Since the seminal paper of Lampson [7] it has been commonplace to view the "protec-

2. There is a certain body of work in distributed logic programming which examines how to transform rule systems in order to minimize communication overhead (e.g. Wolfson and Silberschatz [21]; see also Saraswat *et al.* [22]). However, these works generally do not consider constraints pertaining to boundaries of authority/ownership and privacy of the locally owned rules, that is, limitations as to which processors can be trusted for what.

FIGURE 5. CallerID Example: 'A calling B' Leads to Complex Negotiation.

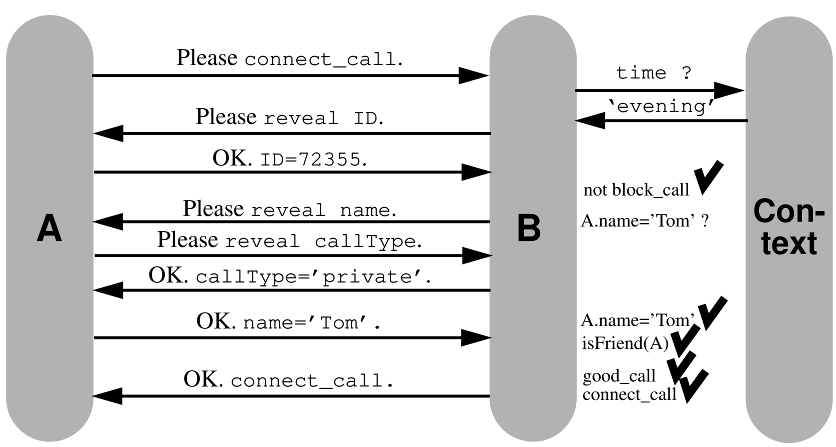

tion" problem as a large global access control matrix (cf. Figure 6), where the human-organizational entities ("subjects"; matrix rows) stand in some authorization relation ("rights"; matrix entries) with information entities ("objects"; matrix columns). This matrix represents the access control problem abstractly. For convenience, it is also common to additionally have "groups" of people included on the subject axis, and, similarly, to have collections of objects on the object axis; such groupings are defined by giving entities "properties."

FIGURE 6. Lampson Access Control Matrix.

	O1	O2	O3	O4	... Objects
S1	r				
S2				r	
S3	r		w		
⋮					
Subjects					

It is then usually noted that this global matrix is impractical to implement directly, and that there are two logical ways of realizing the abstract formulation, which then correspond to the two fundamental conceptualizations which have been investigated in much detail over the past 25 years:

1. **(by column)** ***Access Control Lists (ACLs)*****:** For each object, specify which subjects have which access rights to it.
2. **(by row)** ***Capabilities*****:** For each subject, specify which objects it can access with which rights.

The choice between ACLs and Capabilities is a choice between which one of the two entities, that is, subjects or objects, the access information will be associated with (thus explaining the two names). A combination of the two, a "lock-key" mechanisms, is often used in practice: At first, ACLs are used to determine rights; then these rights are associated as a capability with the corresponding subject.

The relative merits of ACLs versus Capabilities have been investigated and discussed at length (cf. generally a textbook like Silberschatz *et al.* [10]). At this point, we would only like to look into some of the *implicit assumptions* behind the Lampson matrix. In particular, we identify the following assumptions:

- *Subject-Object World:* The basic conceptualization is that we assume a world with a notion of "subject" and "object," that is to say, for instance, that the quality of "subjecthood" comes into being uniformly and independently of the actual interaction. We detail some ramifications of this below (cf. also Thomas and Sandhu [12]). Note also that what is considered "object" here is of course really something provided by another subject, the "owner," that is, the real person which is liable and responsible for it. Indeed, a communication-based model would place this owner right on the same level as the requesting subject. Curiously, in the conceptualization of the access control matrix, owners appear only quite indirectly.
- *Interaction-Independent Objects:* Note that an "object" might come into existence only as part of an interaction. For example, "cgi-bin scripts" of Web servers can synthesize any number of objects at interaction time, without the stipulation that they necessarily "exist" prior to this interaction. We do not doubt that it is possible to conceive a mathematical matrix which covers all these objects, even if this might stretch the idea of a (finite) matrix somewhat. What we want to point to is that this abstraction does not fully reflect the underlying real-world dynamics, and it should not be surprising then if corresponding models end up not capturing certain cases very well.
- *Interaction-Independent Subjects:* A similar thing holds for subjects. Clearly, we know that there are people and groups of people in the world. However, note that the qualities which makes them "subject" or "group" in a given context is assumed here to be ontologically prior to the interaction between "subject" and "object." This "object-property model" assumption is generally far from clear, of course. In particular, it refers to the following assumption.
- *Open-Cards Assumption:* This is the assumption that at the point of the access-control decision what is critical for the decision is laid out "on the table." This corresponds to the case where every person always says everything relevant up-front (only because there might be some property in this description which might qualify for some additional group membership which would then in turn make a difference in the access decision). It is questionable whether this is realistic to assume. Privacy/political considerations are only one such reason: It is the very essence of privacy that persons do not disclose everything "up-front," and that revelation of attributes is modulated in a fine-grained way over time with respect to the present context.

In particular, the trouble of high negotiation complexity in general access control seems to be a result of the assumption on uniform subjecthood in the Lampson matrix, which is not realized in the underlying real-world dynamics. Again, we point this out not to challenge the matrix' mathematical validity and usefulness, but to understand sources of potential problems.

Note that it is generally the case that abstractions which are good but do not fully grasp the underlying dynamics often create artificial "exceptions," that is, certain cases end up not fitting smoothly into the model. It seems that the notion of a “role” of a person is one example of such artificial exceptions, which arises from the fact that subjecthood is treated as ontologically prior to the interaction by which it might only come into being. “Roles” are then invented to try to fix this problem by discretizing subjecthood. (Cf. generally also Winograd and Flores [33].)

5 Communication-based Relationship Model of Access/Action Control

The object-property model of the previous section is not an uncommon conceptualization; we might also call it the “grocery-store model” because it is usually employed as the user-conceptual model, say, when buying an apple for the $1 which it is labelled with. In other words, this conceptualization sees properties such as the price or the access conditions of an object as attributes of the object.

Clearly, for simple exchange interactions of such kind, this is an efficient model, and there is no essential motivation for “contracts” etc. In fact, it has been the case that in underdeveloped economies based preeminently on traditional barter any kinds of exceptions in such “spot market” interactions are quite effectively handled simply by property law claims only (and not by any kind of contract system); indeed it is known that in such environments contract law never developed significantly [19][20]. It is not surprising then that current computing systems also embody this model—the objects dealt with in this context when they were designed certainly resembled more an apple in a grocery store than a piece of land available as “residential real-estate” under certain rental conditions and subject to state and county restrictions about which trees on it its “tenant” (or its “owner”) can cut in which way.

This brings us to the other model of conceiving of “properties,” the contract model, where we think of a subject-subject world with subjects entering into various contractual relations with each other. Properties like price (strictly speaking even ownership, etc.) are then not “in” objects; they are just social, economic, and legal fictions constructed between people on top of objects. Such socially coordinated fictions are then employed to articulate the boundary conditions of people’s interactions in “relational contracts,” which frame relationships, give it structure, and set common expectations.

As we are moving towards richer computing environments that reflect more of the facets of the social sphere, we also have to deal with relations of higher complexities, more subtly constrained sets of rights and obligations, and similar mechanisms for social reference creation, etc. that go beyond the boundary conditions of the technical framework of earlier “protection” designs. For example, software licensing, copyright rules, and usage constraints on personal information certainly begin to resemble in character more a piece of land than an apple in a grocery store.

FIGURE 7. Object-Property Model vs. Agreement Model

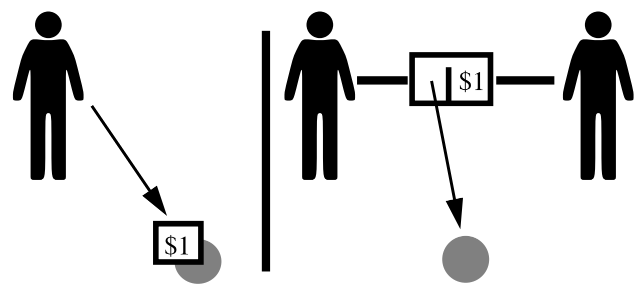

We are interested in this work in a framework designed for dealing with such relationship-based interactions, for which generally only idiosyncratic solutions exist so far; this includes various forms of licensing (e.g. group licensing [4], site licensing, subscription), copyright rights management [24][28][25], privacy, and other more subtle policy issues. In other words, we are not primarily interested in designing for the kinds of spot market transactions at which, for instance, electronic cash is targeted; such mechanisms tie into the agreement model at the level of the actions, though, as specific ways to live up to certain (payment) obligations.[3]

Note that we can see the object-property model as a special case of the communication agreement model for simple exchange cases where the provider is (fairly) indiscriminate about the identity of the other party (“Anyone gets the apple as long as they pay the money for it.”). Care has been taken in the protocol design that such special cases do not generate unreasonable overhead in the agreement model.

Agreement Model of Access/Action Control

In the communication agreement model, we choose a conceptualization which is

- based on a *communication model* of subjects negotiating access and usage conditions with other subjects (the owners of objects if we have objects),[4]

3. Note that the e-cash model is targeted at a spot market environment as it existed in the past, not necessarily as it will be practiced most significantly in the future. In particular, as part of the “universal rise of relationship marketing” (McKenna [14]; cf. also Peppers *et al.* [15]) even spot market transactions such as retail purchases are being recast into a (customer) relationship-centered view with the help of computing and communication technologies to keep track of people individually. In other words, such development only makes more interactions directly relevant under a relationship-based control framework.

4. Note that while this is a shift in perspective for object-metaphor tasks such as access to files or access to digital library contents, it is of course the case that forms of such a subject-subject model have been employed in environments which are themselves already based on communication, for example, (at a much lower level) as part of the communication between LAN routers [13] (albeit not with agreements as first-class citizens and not user-conceptually).

- shifts the perspective from the participants to their *relationship*, and
- introduces an explicit third entity as a social reference point, which encapsulates the relationship in a contract-like object and provides the context for actions (which are then always conducted with respect to it).

We call this entity "commpact," referring to the fact that the relationship object constitutes a compact "communication (com)pact" between the communicating parties here, which explicitly articulates the boundary conditions of their relationship. The notion of a commpact is not supposed to be limited to agreements on the level of legal contracts only; the intention is to extend to a full range of informal usages as well, enabling a shared language for coordinating expectations, in particular also regarding privacy issues. It is this commpact as a first-class citizen around which we center a framework for access/interaction control.

Let us look into a number of points which seem to be essential for understanding what the notion of a contract is about.

- Agreements as a set of enforceable promises can deal with relationships of longer-term nature; they encapsulate boundary conditions of such relationships, and they create social reference points to which people can refer back to at any later point to call into presence what they had coordinated themselves about.
- Agreements can be about objects, but they also uniformly extend to purely relational forms which are not about any objects. Agreements can easily "quantify" over multiple objects. For example, a subscription agreement can be about a whole series of items; there needs to be only one such agreement pointing to the objects about which it is. The same could be achieved in the object-property model only via extensive replication (that is, mentioning all of the conditions in every object). Indeed, agreements can express constraints about objects which do not even exist (yet). For example, subscription agreements are usually about issues which still need to come into existence; nevertheless, we can already talk about these rights and obligations of future objects, pay for them, etc.
- Agreements provide a uniform way for adding a whole number of reservations and special clauses (warranties, guarantees, terms and conditions, etc.); this includes various forms of "strings attached" such as usage conditions. Note that the conventional subject-object model created a gulf which led to the need to separate out "access" and other action (e.g. usage) control. The relationship-based model lends itself quite naturally to uniformly extend to covering usage control issues and obligations and liabilities next to access rights. The agreement serves as a reference point for social coordination about what the intended usage for released information is. The relationship itself is the unit of modulation based on positive or negative feedback.
- Agreements are at least in principle peer-to-peer, not supplicant-granter. The conceptual shift towards centering access/action control around relationships and towards a communication model instead of the supplicant-granter metaphor rephrases the old "access-control" question of "Do I grant this ?" to the new question of "Based on which relationship are we talking to each other ?". It recasts the access control question from that of a (unilateral) "decision" to a matter of agreeing on boundary conditions of a relationship.

Commpact Cards, State, and Personae

In the commpact model, every action is always conducted with respect to some relationship defined in a commpact. The common case is then that the participants *are already in* a relationship. This commpact context can be made present at any point by designating it with an appropriate designator, a "commpact cards," as we call it in analogy to the library card, the student card, the credit card, etc. we have as a token which stands for a relationship (after authentication).

Commpacts are stateful. The simplest case of this is whether or not a commpact is effective between two parties. More complex cases of state are the number of times a license has been used, or whether required notifications have been issued.

"Roles" on the other hand are dealt with by grouping such cards into "personae." For example, people might want to set up a "no-show"/faceless personae, which does not reveal any personal information attributes, uses anonymous payment means, etc.; this personae would group all those commpact cards together which correspond to commpacts that do no reveal any such information. Personae are one example of the kinds of priority lists one might want to maintain on the client-side.

Commpacts in the Lampson Model

As noted above, ACLs and capabilities have been considered as the two fundamental types in which the abstract Lampson matrix can be realized. ACLs associate control information with "objects," capabilities with "subjects" (cf. Figure 8). Note that we can see commpacts as the third logical possibility of realizing the Lampson matrix:

3. (by right type) ***Commpacts*:** For each type of access right, specify which subjects can access which objects.

Or to rephrase it more directly geared towards the intended communication model with objects and access control as an ancillary of this subject-subject relationship:

3.' (by relationship type) ***Commpacts*:** For each relationship type, specify which subjects have an agreement with which object owners about their objects.

Also note that from the point of view of an architecture like the ISO access control framework, the relationship-based model would suggest how to structure the access-control

FIGURE 8.
Realizations of the Lampson Matrix: ACLs, Capabilities, ... and Commpacts.

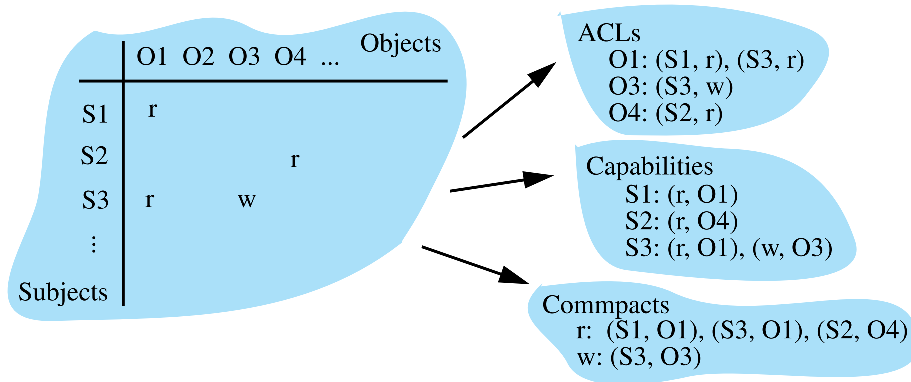

decision module: it would be distributed according to the generic types of relationships two parties can enter, with the ability for users to designate the context of an action. The claim would then be that this structure makes the access control interactions more manageable by taking relationships as the primary.

6 Commpacts: Constraints on Actions

Commpacts are based on a language in which the boundary conditions of a relationship between two persons are expressed by articulating the mutual expectations on rights and obligations. They put constraints on the kinds of actions expected within a certain relationship. When nothing is specified, this can technically be any (inter-)action. Once more rights and obligations are specified, this action space is constrained more narrowly for a given relationship. Two individuals can be in different relationships with each other at a time; certain actions might then be expected under one relationship but not under another. In this way, the commpact language is generalizing such concepts as the "activators" used to express rights and obligations in Minksy [16][17]; it also can be seen as a framework for expressing specific rights systems such as those used in various copyright protection systems [28][25][37][27] or the message types used in electronic contracting systems such as EDI.

Formalism

To write out the commpact language, a basic formalism is needed. The requirements for such a formalism are basically to have the ability to express nested attribute-value structures plus a way to indicate sharing of substructures and a way to assert constraints on values. We have been using for this purpose Attribute-Value Matrices (AVMs), a mathematically well-understood unification-based constraint language which has proven useful in linguistics to describe constraints on utterance-semantics relationships [23]. Unification is a well-defined way to deal with partiality. Since we are often interested in assembling partial descriptions from different authorities, unification serves as a natural mechanism by which composite structures can be put together from partial, possibly redundant, distributed structures. An associated simple constraint language is used to express constraints between values.

Commpacts, E-persons, and Credentials

Commpacts are considered to be between "e-persons"/ "epers", a legally motivated notion introduced by Karnow [8] to provide a shield for privacy much in the same way as, for instance, the notion of a corporate person provides a shield which safeguards personal assets from business failures or work-related lawsuits. Such epers are identified by a *minimal principal handle* [29], which is also the unit by which persons are authenticated in the absence of more detailed attributes (dealt with then via credentials).

Commpacts are technically a set of enforceable promises:

```
Commpact {
  name      Name
  purpose   Description
  state     State
  partyA    Eperson
  partyB    Eperson
  precCond  PrecedentCond
  promises  SET OF Promise
  secSpec   SecuritySpec
  objects   ObjectHandles
  tncs      TermsAndConditions
  }
```

Credentials play an important part in commpacts; they are conceived here generally as certification of arbitrary properties by different authorities [29]. For example, a subscription commpact for students with a certain minimum age might require credentials from appropriate authorities about age and college affilitation.

Rights and Obligations

Rights include but are not limited to rights about (owned) objects, that is, for instance, the conventional intellectual property rights [27][37]. Obligations are promises for actions bound to happen in the future. An example would be the obligation to have established a prerequisite commpact, possibly between different parties. Having to present credentials about personal information attributes is an example of a condition which can either be framed as a precedent condition (conditional contract) or as a promissory condition as part of an obligation.

```
Promise {
 name       Name
 descrip    Description
 holder     Eperson
 promCond   PromissoryCond
 condSub    ConditionSubsequent
 effective  TimePlace
 enfSpec    EnforcementSpec
 }

Right:Promise {
 exclusive  Boolean
 canDo      Action
 }

Obligation:Promise {
 mustDo     Action -- covers "refrain from"
 otherwise  Action
 }
```

Actions are further constrained at the level of primitive actions which can or have to take place; these include a version of the intellectual property actions of the copyright act (reinstantiating and using), commpact formation actions (accept, terminate, etc.), notification, as well as a range of domain-specific actions such as payment, etc.

Example: Subscription Commpact

A typical example of a commpact would be a subscription agreement between a consumer and a newsletter publisher. Such a subscription commpact typically would be derived from a standard document access commpact; it would have a general section which describes such attributes as name, purpose, security level, persistence requirements (to which extent state is kept persistent and for how long), trust levels (which authorities are trusted for what) as well as what is defined as constituting state. A more rights-specific section would then detail qualifying requirements (which prerequisites each party has to fulfill in order for there to be a valid offer or commpact) as well as identification requirements (identifier and which additional credentials are required).

Then, the rights and obligations from this agreement would be listed and linked to the corresponding attribute values and constraints. For example, a sample commpact might specify here the usage rights for the newsletter issues, that is, which issues can be accessed, whether they can be copied and redistributed (and according to which other commpact then); it might also specify a payment schedule and the usage rights for the personal information which the subscriber makes available. The subscriber could then choose in his client-side (commpact interpreting) interface how to fulfill his obligations. One possibility might be automatic monthly payment; another a simple notification at the beginning of each quarter (and perhaps off-line payment).

7 Interfaces and Protocols

This section briefly surveys the basic object-request interfaces and the protocols.

Object-Request Interfaces

Four basic separate out functionality for different usages:

- *Commpact Form Definition Interface:* This interface allows for entirely new forms of commpacts to be defined. We expect this to be an infrequent case**;** when it does happen, it would be generally dealt with by professionals at the trusted proxies such as home provider administrators. This interface is analogous to coming up with a new standard rental agreement form, that is, a task which most end-users would generally not take on themselves; at best they would want to customize an existing form. Commpact forms give such flexibility by allowing people to insert different certificate types and instantiate a commpact's parameters differently. Note that this interface corresponds to the one in EDI of setting up a new (UNSM) "EDI standard message type" (which involves submitting a request to the RT secretariat of the UN/ECE/WP.4).
- *Commpact Negotiation Interface:* This is for the case where two participants want to set up a new relationship based on standard commpact forms. Cf. below.
- *Action Request Interface:* This is what the design is optimized for as the common used one; an action is supposed to be performed based on an already established relationship. In most cases, this can happen by just referencing the commpact context with a commpact card.
- *Certification Request Interface:* This is the interface for establishing, renewing, or revoking certificates about arbitrary attributes; one way to do this is with credentials at the certification manager's back-end (cf. [27]).

Negotiation

Commpacts are agreed upon as a result of a negotiation according to a general, domain- and content-independent protocol, which is designed to reflect legal contract practices [3][35]. Figure 9 shows a finite-state diagram which defines the sequences in a negotiation process that are leading to a successful commpact formation. We expect this negotiation to be generally quite succinct, more similar to the EDI message exchange than to an auctioning process.

FIGURE 9. Commpact Negotiation: Sequencing.

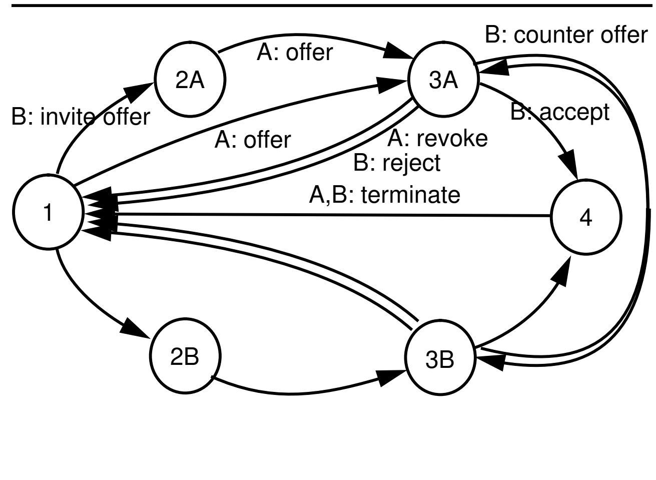

FIGURE 11. CallerID Example: 'A calling B' with Commpacts in Trusted Proxy Environment.

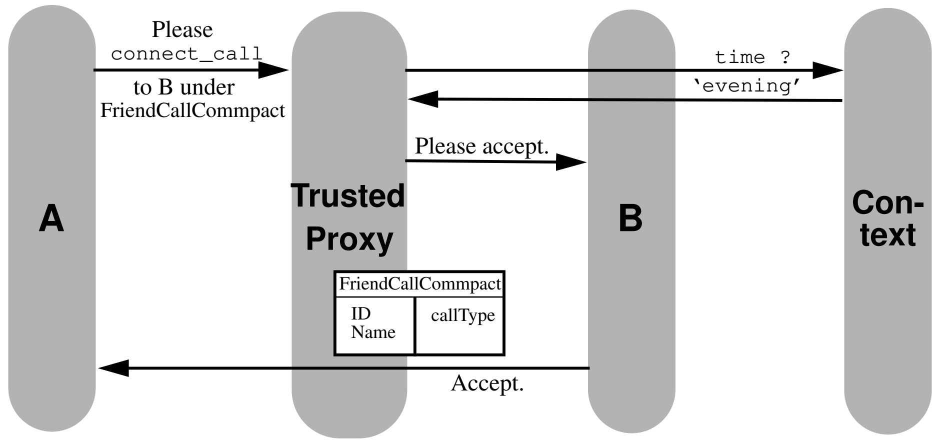

Note that appropriate initial commpact forms can be found in the simplest case by some form of browsing.

Commpact Replication

Trusted proxies are exploited in the commpact framework in that the initial commpact forms will usually reside with them. Once initiated, there is an orthogonal choice of replicating a commpact (typically at the client- and/or the server-side). This corresponds to the real-world case where an agreement is distributed in copies to the contract parties instead of being deposited with a third party. Replication plays a role in enforcement issues, which is being dealt with on multiple levels in the commpact framework, the most secure version being based on the standard trusted computing base concept.

FIGURE 10. Orthogonal Choice: Commpact Replication.

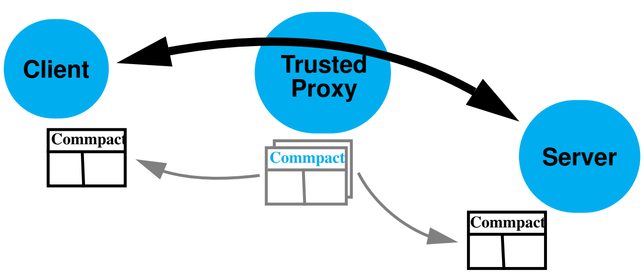

8 CallerID Example Revisited

Let us look here how the "CallerID" interactions of Figure 5 would look like in a communication agreement framework (cf. Figure 11). The commpact framework would essentially support people to first agree on which relationship to establish with which other persons. This includes commpacts with phone marketers corresponding to the rules of Figure 4. A caller would then have to choose based on which commpact to make a call; the default for a telemarketer would then of course be an appropriate telemarketer commpact unless the agreement will be violated.

Note that the trusted proxy knows about which expectations B has towards A, and which ones A has towards B; it can thus easily coordinate the mutual expectations. If B accepts (possibly based on another decision facility), then the connection will be established. In other words, commpacts at a trusted proxy substitute proxy-internal processing for negotiation via the network, and the alignment into commpacts serves as a form of preprocessing which can be performed once the rules have been partially localized in a trusted proxy.

A number of special cases which have been raised as objections against the introduction of Calling-Number identification in phones can be readily dealt with. For example, it was pointed out in the debates surrounding the CallerID issue that certain people like psychiatrists might want to call others (patients) without revealing their number because they would allegedly run certain risks then. This points to the necessity of a blocking feature for a certain set of circumstances. On the other hand, if an individual calls 911 in case of emergency, then the blocking feature, which might be enabled under certain circumstances, should be inactivated since otherwise the caller cannot be located. In the commpact model, each of these types of behaviors would basically get a different commpact. Next to the commpacts mentioned above, there might simply be an `EmergencyCallCommpact` for calls to 911, and an `PatientCallCommpact` for cases such as the one mentioned above. They would deal with the special cases without interfering with any of the other rules implicitly articulated in any of the other commpacts.

In other words, the encapsulation of a commpact provides a framework for dealing with the various exceptions in a uniform way. By having commpacts as social reference points one has a way of talking about expectations from a certain relationship, and one has a base-line against which to measure contract performance. For example, a telemarketer using a `FriendsCallCommpact` for professional calls would violate such expectations; if this was not already ruled out by technical enforcement, the reputation-based enforcement of the relevant home providers could then deal with such cases. Note that this widens up the action-interrupt enforcement paradigm of the ISO model to a more flexible range of enforcement options (detailed further in [30]).

9 Conclusion

We have introduced a relationship model of access control which is based on a model in which peers explicitly negotiate the boundary conditions of their relationship in a user-designatable communication pact.

The notion of an agreement between a requester and the owner of requested services and objects provides a uniform user-conceptual model for access control issues (and associated rights and obligations) in a complex social, economic, and legal environment. It provides an encapsulation which manages complexity in a way that can translate into usability benefits while at the same time not curtailing the ability of distributed parties to contribute their preferred access policies in a modular way.

The framework maps naturally onto a trusted proxy implementation model, which we argue is the environment for which there is a need for user-conceptually uniform access control models that can deal with the rich usages of emerging networked environments. Moreover, the framework is designed to leverage the known versatility of reputation-based enforcement institutions by creating social reference points and by providing a language for a mechanism of coordinated expectation.

Acknowledgements

Our sincere thanks go to all people who provided feedback.

This paper is based upon work supported by the National Science Foundation under Cooperative Agreement IRI-9411306. Funding for this cooperative agreement is also provided by ARPA, NASA, and the industrial partners of the Stanford Integrated Digital Libraries Project. Any opinions, finding, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation or the other sponsors.

10 References

[1] Biba (1977). Integrity Considerations for Secure Computer Systems. ESD-TR-76-372, Electronic Systems Division, Bedford, MA.

[2] Clark, D.D., and D.R. Wilson (1987). A Comparison of Commercial and Military Security Policies. *IEEE Symp. on Security and Privacy.*

[3] Eisenberg, M.A. (1985). *Contracts. Gilbert Law Summary.* Harcourt, Brace, Jovaovich Legal and Professional Publications.

[4] Hauser, R. (1993). Does Licensing Require New Access Control Techniques ? *Proceedings of the First ACM Conference on Computer and Communications Security*, pp. 1-8. Fairfax, VA.

[5] La Padula, L. (1990). Formal Modeling in a Generalized Framework for Access Control. Proc. *IEEE Symposium on Security and Privacy.*

[6] ISO (1989). Security Framework III: Access Control Framework. ISO/IEC JTC1/SC21 N4206. Draft, November.

[7] Lampson, B.W. (1971). Protection. *5th Princeton Symposium on Information Science and Systems.* Reprinted in *ACM Operating Systems Review* **8**(1):18-24, 1974.

[8] Karnow, Curtis E.A. (1994). The Encrypted Self: Fleshing out the Rights of Electronic Personalities. *Conference on Computers, Freedom, and Privacy.*

[9] Nelson, C. (1995). The ABC of EDI. *EDI Aware*, Issue 4, Winter. Cf. URL http://infopole1.soca.cf.ac.uk/edi/EDIAware4Index.html.

[10] Silberschatz, A., J. Peterson, and P.G. Galvin (1991). *Operating Systems Concepts*. Addison-Wesley.

[11] Sandhu, R.S. (1992). The Typed Access Matrix Model. *Proceedings of the IEEE Symposium on Security and Privacy.*

[12] Thomas, R.K., and R.S. Sandhu (1994). Conceptual Foundations for a Model of Task-based Authorizations. In *Proceedings of the IEEE Symposium on Security and Privacy.*

[13] McHugh, J. and A. P. Moore (1986). A Security Policy and Formal Top-level Specification for a Multi-level Secure Local Area Network. In *IEEE Symposium on Security and Privacy*, Oakland.

[14] McKenna, R. (1991). *Relationship Marketing: Successful Strategies for the Age of the Customer.* Addison-Wesley.

[15] Peppers, D., and M. Rogers (1993). *The One to One Future: Building Relationships One Customer At a Time.* Doubleday.

[16] Minsky, N. (1978). An Operation-Control Scheme for Authorization in Computer Systems. *International Journal of Computer and Information Sciences* **7**(2), pp.157-91.

[17] Minsky, N.H., and A.D. Lockman (1985). Ensuring Integrity by Adding Obligations to Privileges. *Proceedings of the 8th International Conference on Software Engineering*, pp. 92-102.

[18] Moffett, J.D., and M. S. Sloman (1991). Content-dependent Access Control. *Operating Systems Review* **25** (2), pp. 63-70, April.

[19] Milgrom, P., and J. Roberts (1992). *Economics, Organization, and Management*. Prentice Hall, NJ.

[20] Coase, R.H. (1990). *Institutions, Institutional Change and Economic Performance.* Cambridge University Press.

[21] Wolfson, O., and A. Silberschatz (1988). Distributed Processing of Logic Programs. *ACM SIGMOD International Conference on Management of Data* 1**7**(3), pp. 329-36.

[22] Saraswat, V., K. Kahn, and J. Levy (1990). Janus: A Step towards Distributed Constraint Programming. *North American Conference on Logic Programming*. Austin, TX. pp. 431-46. MIT Press.

[23] Sag, I., and D. Pollard (1991). HPSG. CSLI Publications.

[24] EPR (1995). NetTrust Electronic Rights System. Electronic Publishing Resources, Incorporated. URL: http://www.epr.com/.

[25] Stefik, M. (1995). Letting loose the light: Igniting commerce in electronic publishing. Report, Xerox Palo Alto Research Center.

[26] Ullman, J.D. (1989). *Database and Knowledge-base Systems: The New Technologies*. Volume II. Computer Science Press.

[27] Perritt, H. (1994). Permission Headers and Contract Law. *Proceedings of the Workshop on Technological Strategies for Protecting Intellectual Property in the Networked Multimedia Environment*. Coalition for Networked Information.

[28] Kahn, R.E. (1994). Deposit, Registration, and Recordation in an Electronic Copyright Management System. In *Proceedings of Technical Strategies for Protecting Intellectual Property in the Networked Multimedia Environment*, Annapolis, MD.

[29] Röscheisen, M. (1995). General Certificates. Working Paper #12. Stanford Integrated Digital Libraries Project, Stanford University.

[30] Röscheisen, M. (1996). Panoptic Access/Action Control. Working Paper. Computer Science Department, Stanford University.

[31] Röscheisen, M. (in progress). *Content and Access Control including Privacy: Design for Relationships.* Dissertation. CS Dept., Stanford University. URL: http://diglib.stanford.edu/rmr/thesis/.

[32] Winograd, T. (1995). The Proxy Is Where It's At. Working Paper #8. Stanford Integrated Digital Libraries Project, Stanford University.

[33] Winograd, T., and F. Flores (1986). *Understanding Computers and Cognition: A New Foundation for Design*. Addison-Wesley.

[34] Greguras, F.M., T.A. Golobic, R.A. Mesa, R. Duncan (1995). Online Contract Issues. Updated version of a presentation made at *Law Seminars International Electronic Commerce: Doing Business Online*, September 21, 1995.

[35] Wright, B. (1995). *The Law of Electronic Commerce. EDI, E-Mail, and Internet: Technology, Proof, and Liability.* Little, Brown & Co.

[36] Saltzer, J.D., and M.D. Schroeder (1975). The Protection of Information in Computer Systems. *Proc. of IEEE* **63**(9), pp. 1278-1308.

[37] Upthegrove, Luella, and T. Roberts (1994). Intellectual Property Header Descriptors: A Dynamic Approach. *Proc. of the Workshop on Technological Strategies for Protecting Intellectual Property in the Networked Multimedia Environmentt*, Annapolis, MD.