Report Number: WUCSE-2003-70

2003-10-14

Managing Access Control in the Presence of Mobility

Christine Julien, Gruia-Catalin Roman, and Jamie Payton

The increased pervasiveness of wireless mobile computing devices draws new attention to the need for coordination among small, networked components. The very nature of the environment requires devices to interact even when they meet unpredictably. Because these networks are often decoupled from a fixed infrastructure, reliance on centralized servers for authentication and access policies is impractical. Access control is crucial in such systems, and applications must directly manipulate and examine access policies because they require full control of their data. In this paper, we explore the essential features of general access control policies tailored to the needs of agent... Read complete abstract on page 2.

Managing Access Control in the Presence of Mobility

Christine Julien, Gruia-Catalin Roman, and Jamie Payton

Complete Abstract:

The increased pervasiveness of wireless mobile computing devices draws new attention to the need for coordination among small, networked components. The very nature of the environment requires devices to interact even when they meet unpredictably. Because these networks are often decoupled from a fixed infrastructure, reliance on centralized servers for authentication and access policies is impractical. Access control is crucial in such systems, and applications must directly manipulate and examine access policies because they require full control of their data. In this paper, we explore the essential features of general access control policies tailored to the needs of agent coordination in the presence of physical and logical mobility. We propose and evaluate novel constructs to support such policies in mobile applications.

Department of Computer Science & Engineering

Washington University in St. Louis
School of Engineering & Applied Science

2003-70

Managing Access Control in the Presence of Mobility **PLEASE SEE WUCSE-04-11**

Authors: Julien, Christine; Roman, Gruia-Catalin; Payton, Jamie

October 14, 2003

Abstract: The increased pervasiveness of wireless mobile computing devices draws new attention to the need for coordination among small, networked components. The very nature of the environment requires devices to interact even when they meet unpredictably. Because these networks are often decoupled from a fixed infrastructure, reliance on centralized servers for authentication and access policies is impractical. Access control is crucial in such systems, and applications must directly manipulate and examine access policies because they require full control of their data. In this paper, we explore the essential features of general access control policies tailored to the needs of agent coordination in the presence of physical and logical mobility. We propose and evaluate novel constructs to support such policies in mobile applications.

Managing Access Control in the Presence of Mobility

Christine Julien, Gruia-Catalin Roman, and Jamie Payton
Department of Computer Science and Engineering
Washington University in St. Louis
{julien, roman, payton}@wustl.edu

Abstract

The increased pervasiveness of wireless mobile computing devices draws new attention to the need for coordination among small, networked components. The very nature of the environment requires devices to interact even when they meet unpredictably. Because these networks are often decoupled from a fixed infrastructure, reliance on centralized servers for authentication and access policies is impractical. Access control is crucial in such systems, and applications must directly manipulate and examine access policies because they require full control of their data. In this paper, we explore the essential features of general access control policies tailored to the needs of agent coordination in the presence of physical and logical mobility. We propose and evaluate novel constructs to support such policies in mobile applications.

1. Introduction

Ubiquitous computing devices communicate wirelessly, opportunistically forming ad hoc networks not connected to a wired infrastructure. In such environments, distributed applications exchange information or coordinate tasks. These networks can include a handful of devices or thousands of heterogeneous components, making coordinating and mediating their competing needs a massive task. Much research focuses on developing middleware to facilitate rapid application development for this demanding environment.

This paper focuses on systems that use tuple spaces. Linda [6] provides a centralized tuple space where application agents exchange information using content-based matching of patterns against data. Variations on this theme adapt it to the mobile environment where a central repository is not feasible. Due to the open and dynamic nature of mobile systems, security concerns of three types arise: protecting mobile hosts from malicious agents, protecting agents from tampering hosts, and securing data. D'Agents [7] uses public-key cryptography to authenticate incoming agents to increase host security. Undetachable threshold signatures [1] prevent hosts from tampering with an agent's data.

Protecting data include both ensuring data integrity and controlling access. Much coordination research has addressed the former by encrypting communication within coordination spaces. SAMCat [12] and Yalta [3] use encryption and authentication to securely transmit tuples into and out of a data space. Our work focuses on the final issue: controlling data access. A solution to this problem is complicated by the fact that, in the mobile environment, disconnection from a wired infrastructure renders a centralized solution impossible.

A common mechanism for addressing access control uses access matrices to describe rights. The rows of the matrix correspond to users and the columns to objects; a cell in the matrix contains the access rights a user has on an object. This approach generalizes several commonly used approaches, including access control lists and capability definitions. In the mobile environment, the number of possible agents and the amount of data available over the lifetime of the system makes applying these solutions directly impractical. The access control function introduced in this paper overcomes the limitations imposed by mobile systems by operating over general descriptions of interacting parties and dynamically adjusting to the changing context.

Section 2 introduces a general mobile coordination model. Section 3 describes our access control mechanism. In Section 4, we discuss the construct's expressive power and overhead. Section 5 overviews related work, and conclusions appear in Section 6.

2. A Generalized Coordination Model

In this section, we capture the essential features of tuple space coordination in order to explain access control requirements for mobile middleware. The result is a generalization that spans the gamut from tuple definition to sophisticated operations. In the original Linda

model, processes generate tuples in a centralized repository and retrieve them using content-based operations in which the retrieving process specifies a pattern that the returned tuple must match. These operations are synchronous in that they "block" the issuing process until a tuple satisfies the operation and is returned. The Linda operations decouple agents in a manner useful in mobile networks, as demonstrated below.

The Tuple Space. Some mobile systems (e.g., MARS [4]) focus on logically mobile agents in a network of physically stationary hosts, while other systems (e.g., LIME [11] and EgoSpaces [9]) integrate physical and logical mobility. All such systems associate a tuple space with a network component that allows other components to access the data. Tuple spaces can be permanently bound to hosts, to agents, or distributed among a combination of the two. The distribution of the tuples is irrelevant with respect to access control; the key aspect of the representation is how application agents access data. We assume a tuple space bound to each mobile agent. Using this model, we can simulate other approaches, e.g., to simulate tuple spaces bound to a host, we associate an agent permanently to a host and use its tuple space as the host's tuple space.

Tuples and Patterns. We generalize a tuple to one in which a field is identified by a name. A tuple is an unordered set of triples: $\langle(name, type, value), \ldots\rangle$. For each field, *type* is the data type of *value*. In a tuple, each field *name* must be unique. Users access tuple spaces by matching patterns against tuples. A pattern has the form: $\langle(name, type, constraint), \ldots\rangle$. The *constraint*s are functions that provide requirements a field's *value* must match for the tuple's field to match the pattern's field. Specifically, the matching function $\mathcal{M}$ is defined over a tuple θ and a pattern p as:

$$\mathcal{M}(\theta, p) \equiv \langle \forall c : c \in p :: \langle \exists f : f \in \theta \wedge f.name = c.name \wedge f.type \texttt{ instanceof } c.type :: c.constraint(f.value)\rangle\rangle.$$ [1]

$\mathcal{M}$ requires that, for every constraint c in the pattern, there is a field f in the tuple with the same name, the same type or a derived type, and a value that satisfies c. While the function requires that each constraint is satisfied, it does not require that every field in the tuple is constrained, i.e., a tuple must contain all the fields in the pattern but can contain additional fields.

Basic Operations. Next, we classify the available operations, regardless of the tuple space structure.

1 In the notation $\langle$**op** *quantified_vars* : *range* :: *exp*$\rangle$, the variables from *quantified_vars* take on all values permitted by *range*. Each instantiation of the variables is substituted in *exp*, producing a multiset of values to which **op** is applied, yielding the value of the three-part expression. If no instantiation of the variables satisfies *range*, the value of the expression is the identity element for **op**, e.g., *true* when **op** is $\forall$.

Tuple Generation. Agents create tuples using **out** operations. Tuple generation generally places a tuple (t) in a specific tuple space: **out**(T, t), where T is a tuple space with a particular name located at a particular agent. In EgoSpaces, an **out** places the tuple in a local tuple space controlled by the generating agent. In LIME an **out** can place a tuple in any tuple space owned by any agent on a connected host. In MARS the tuple is created in the local host's tuple space.

Tuple Retrieval. To read and remove tuples, agents use **rd** and **in** operations respectively, which assume three forms: blocking, atomic probing, and scattered probing. The blocking form, **rd**(T, p), returns a tuple matching the pattern p from the tuple space T. The tuple space can be either local to the agent or controlled by another network component. Atomic probing operations, **rdp** and **inp**, guarantee, if a matching tuple exists, it is returned, but they can return ϵ if no match immediately exists. Like the blocking operations, they are atomic with respect to the tuple space on which they are issued; in some cases in the mobile environment, guaranteeing this atomicity can be expensive. Scattered probing operations, **rdsp** and **insp** offer weaker guarantees. While all of these access operations entail only single tuples, extensions to Linda allow simultaneous access to groups of tuples. These operations come in all three forms described above and are referred to as group operations, e.g., **rdg** refers to a blocking non-destructive read operation that returns all matching tuples from the tuple space.

Some models present tuple space operations to the user in a different manner. In LIME, application agents operate over a federation of connected tuple spaces, while in EgoSpaces, agents operate over projections, called *views*, of all available data. In these cases, the more complex interactions can be reduced to the tuple space operations described above.

3. Access Control Function

As dynamic components become increasingly pervasive, security concerns become of paramount importance. In our coordination model, an agent assumes responsibility for mediating access to its data. The ability to control access in this manner is fundamental because it allows the access policies to reflect an agent's instantaneous needs. To accomplish this, each agent specifies an individualized access control function.

We allow an agent to restrict which agents access its data and the manner in which the access occurs. To accomplish the former, a requesting agent must provide credentials identifying itself. To accomplish the latter, the access control function accounts for the op-

eration being performed. In the end, each agent defines a single access control function that takes as parameters a tuple, a set of credentials identifying the requesting agent, the operation being performed, the pattern used in the operation, and the owning agent's profile (defined next). This function returns a boolean indicating whether the requested access is allowed.

Profiles. Before describing the access control function in more detail, we introduce a profile to maintain properties of each agent, which we represent as a tuple. Particular applications or coordination systems may require specific attributes in this profile. In general, we assume a profile contains at least a unique host id identifying the agent's host and a unique agent id.

Parameters. An access control function takes five parameters: the credentials, operation, tuple, pattern, and the owner's profile.

Credentials. Credentials allow an agent to convey information about itself. In simple cases, they can be a standard set of attributes, e.g., the agent's id or a third-party authentication. When an agent has a priori knowledge of the access requirements, credentials can be more complicated, e.g., a password. When constructing credentials, an agent must take care not to give away too much information, e.g., if the agent has multiple passwords, it should send only the correct one. This identification is especially necessary in open and dynamic mobile environments, where it is often not possible to know a priori exactly which agents can access restricted information. Instead, agents must prove they have required privileges. Credentials are a subset of an agent profile presented as a tuple of attributes, which allows the access control function to use pattern matching to evaluate credentials. The credentials and their transmission with the operation are assumed to be private. This security is outside the scope of this paper but could be accomplished using cryptography schemes already under development.

Operation. The access control function can also account for the operation requested. Often, some data should be restricted to read-only access, yet current systems do not inherently allow this restriction. Considering the operation when determining access allows a dynamic application to permit one set of operations for some agents, but different operations for others.

Requested Tuple. Because we focus on tuple space models, the access control function can operate over the tuple to be returned from an operation. Pattern-matching allows this portion of the access control function to be easily defined while remaining flexible.

Pattern. A powerful component of the access control function is its ability to account for the pattern used in the content-based operation. The pattern provides information about an application's prior knowledge of the data. The owning agent may allow access only to agents that know the "correct" way to access the data (e.g., providing a wild card pattern that matches any tuple may not be acceptable). Some knowledge of the structure of the requested tuple might indicate that the requesting agent shares common application goals.

Owner's Profile. The access control function also considers the owner's current state. Because the access policy is determined dynamically, access can be granted based on context information. In some cases, data may never be sent wirelessly between devices unless they are within a secure physical environment where eavesdropping is known to be impossible.

Access Control Function. The access control function takes the five parameters described above, and determines whether or not to allow the requested access. Formally, this function can be represented as: $\text{ACF} : T \times C \times O \times P \times \Pi \rightarrow \{0,1\}$, where T is the universe of tuples, C is the universe of credentials, O is the finite set of operations, P is the universe of patterns, and Π is the universe of profiles. The access control function (ACF) maps the values of the parameters to a boolean indicating the access decision. The function can also be represented as: $access = \text{ACF}(credentials_r, op, tuple, pattern, profile_o)$; r is the requesting agent and o is the tuple's owner.

We will briefly discuss the expressive power of this construct later. For now we consider what it *cannot* easily represent. Access decisions cannot be based on properties of the requesting agent not included in its credentials. The requesting agent must carefully construct the credentials it sends with each operation request. The decision can also not rely on arbitrary environmental properties. For example, an agent cannot base a decision on the number of copies of a tuple. The access control function lends itself well to the mobile environment because it allows access policies to adapt to the context. Access decisions are transparent to requesting agents; if access is denied, a requester does not even know that the matching tuple existed.

Using the Access Control Function. We first show how the access control function benefits a particular middleware system, EgoSpaces. We then show how restricting operations to administrative domains can be implemented with the new construct.

Use in EgoSpaces. EgoSpaces addresses the needs of agents in large-scale heterogeneous environments. An agent operates over a context that can include, in principle, all data in an entire network. EgoSpaces' unique model of coordination, however, structures data in terms of *views*, or projections of the maximal set of data. Each agent defines its own views; these individu-

alized views abstract the dynamic environment by constraining properties of the network, hosts, agents, and data. To further reduce programming costs, EgoSpaces transparently maintains views; as hosts and agents move, the view's contents automatically reflect context changes without the agent's explicit intervention. EgoSpaces can employ an agent-specified access control function on a per-view basis. When an agent defines a view, it attaches a set of credentials and a list of operations it intends to perform on the view. EgoSpaces uses each contributing agent's access control function to determine which tuples belong in the view. In the end, the view contains only tuples that qualify via their owning agent's access control function.

Administrative Domains. Some applications restrict agent operations to administrative domains. For example, assume nested domains defined as a university's computers, a department's computers, and a research group's computers. To provide security guarantees, applications limit access to certain data to only computers on the university's network. Still other data ought to be restricted to departmental computers, or to research group computers. A user in the research group, working on a mobile computer, wants to use a software license of which the research group has n copies. The licenses are stored as tuples in a tuple space. Each computer in the group carries a tuple space; the available licenses are initially distributed in some random fashion. A user can take a license if it is not in use and the user holding the license is within communication range. The agents controlling the licenses restrict access to only group members who have departmental authentication (retrieved a priori), and are running on computers in the university domain. To retrieve a license, a user provides these three properties as credentials and attempts to **in** a license from a connected tuple space. If successful, the number of available licenses decreases by one. When the user finishes using the software, it replaces the license in its local tuple space.

4. Discussion

The access control function provides a flexible mechanism for agents to specify privileges dynamically and adaptively in mobile coordination models.

Expressiveness. While its expressiveness makes the access control function more flexible and arguably more useful in coordination among constantly changing mobile agents, this flexibility comes with some cost.

Credentials. On one hand, because credentials can encode arbitrary information about an agent, particular applications can adapt credentials to their needs. On the other hand, a requesting agent must not reveal too much information since any information sent in credentials is no longer secret.

Functions. Because the access control function takes a number of parameters, an agent can dynamically adjust its access policies. Again, flexibility comes with a cost. While complex access control policies are possible, constructing the function (from the developer's perspective) can become difficult. Fortunately, the design of the function prevents this complexity from affecting agents that do not require complex policies.

Overhead. Given the model's expressiveness, it is useful to evaluate its overhead. The addition of the access control mechanism introduces some amount of programming overhead. This overhead is difficult to quantify without a case study involving users implementing actual access control policies. While this is a useful future task, it is outside the scope of this paper.

Additional Communication. The key aspect of the communication overhead is the amount of data (in bits) that must be sent. Before adding the access control mechanism, the number of bits required to send an operation request is: $b = |op| + |pattern| + |agent\_id_r|$, where $|op|$ is the number of bits required to identify the operation. $|pattern|$ is the number of bits required to represent the pattern. This depends on the number of fields in the pattern. $|agent\_id|$ is the number of bits required to identify the requesting agent so the response can be returned. It is likely that the pattern, which encodes the content-based nature of the request, dominates this expression, as the op and $agent\_id_r$ are simple data types with small, constant lengths.

We can write a similar term to express the number of bits needed to be sent when using the access control function. This includes only the addition of the number of bits necessary to encode the credentials: $b_{acf} = |op| + |pattern| + |agent\_id_r| + |credentials_r|$.

Credentials are a tuple. Because tuples are similar to patterns the number of bits required to represent the credentials is likely near the number of bits needed to represent a pattern. If so, the overhead of using access control is approximately 2. An application can directly control the amount of overhead it incurs because it determines what credentials to send with each request. In this respect, the use of application intuition to reduce the credentials transmitted to exactly those required reduces the overhead of the communication.

Additional Computation. Evaluating the access control function also requires additional computation in the form of an additional method invocation. Because the function can contain arbitrary code, the computational overhead lies in the hands of the application programmer. From the programmer's perspective, the operating conditions of the application must be a primary

concern. If so desired, a system can include a mechanism to prevent undesirable access control functions by bounding the time they are allowed to run or by imposing restrictions on their capabilities.

5. Related Work

As discussed previously, the access matrix does not directly lend itself to mobile systems. In one example of attempting to apply such a method, TuCSoN agents [5] are assigned capabilities defining tuple space operations for particular patterns in a certain tuple space. An access control list for the tuple space stores these capabilities. This approach requires that all coordinating parties are known in advance and that a centralized party determines access policies statically.

Other systems use encryption for access control. In SecOS [2], tuples are unordered sequences of individually encrypted fields, and, to match an encrypted field, a pattern must contain a correct key. Other work [8] associates keys with tuple spaces, and an agent must provide the key to access the tuple space. While both of these models provide access control mechanisms, they require secure key distribution and management, which affects the scalability of the system.

Law Governed Interaction (LGI) [10] provides an expressive approach to access control in which agents must adhere to a law that imposes context-sensitive constraints on the execution of tuple space operations. A law dictates actions an agent performs in response to the arrival or departure of tuple space operations. Programming applications in LGI requires programming specific actions in the access control policy and adding a controller to mediate tuple space requests. In contrast, in our model, programming takes place in the coordination model, and the agent's requested operation is checked with the access control function.

6. Conclusion

In this paper, we first provided a generalized coordination model representative of those used in dynamic pervasive computing environments. We then introduced access control functions for mobile coordination and showed how they could be successfully used in these systems. While this construct does incur some overhead, the expense is not prohibitive when compared with the benefits it offers. The novel access control function directly addresses the specific access control needs of mobile coordination models. In particular, the construct provides increased scalability and decoupling when compared with previous approaches, without sacrificing flexibility and expressiveness.

ACKNOWLEDGEMENTS

This research was supported in part by the Office of Naval Research under ONR MURI research contract N00014-02-1-0715. Any opinions, findings, and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the Office of Naval Research.

References

[1] N. Borselius, C. J. Mitchell, and A. Wilson. Undetachable threshold signatures. In *Cryptography and Coding—Proc. of the 8th IMA Int'l. Conf.*, volume 2360 of *LNCS*, pages 239–244, 2001.

[2] C. Bryce, M. Oriol, and J. Vitek. A coordination model for agents based on secure spaces. In P. Ciancarini and A. Wolf, editors, *Proc. of the 3rd Int'l. Conf. on Coordination Models and Languages*, pages 4–20. Springer-Verlag, 1999.

[3] G. Byrd, F. Gong, C. Sargor, and T. Smith. Yalta: A secure collaborative space for dynamic coalitions. In *IEEE 2nd SMC Info. Assurance Workshop*, 2001.

[4] G. Cabri, L. Leonardi, and F. Zambonelli. MARS: A programmable coordination architecture for mobile agents. *Internet Computing*, 4(4):26–35, 2000.

[5] M. Cremonini, A. Omicini, and F. Zambonelli. Coordination and access control in open distributed agent systems: the TuCSoN approach. In A. Porto and G.-C. Roman, editors, *Coordination Languages and Models*, volume 1906 of *LNCS*, pages 99–114. Springer-Verlag, 2000.

[6] D. Gelernter. Generative communication in Linda. *ACM Transactions on Programming Languages and Systems*, 7(1):80–112, 1985.

[7] R. Gray, D. Kotz, G. Cybenko, and D. Rus. D'Agents: Security in a multiple-language, mobile-agent system. In G. Vigna, editor, *Mobile Agents and Security*, volume 1419 of *LNCS*, pages 154–187. Springer-Verlag, 1998.

[8] R. Handorean and G.-C. Roman. Secure servise provision in ad hoc networks. In *Proceedings of the 1st Int'l Conf. on Service Oriented Computing.* (to appear).

[9] C. Julien and G.-C. Roman. Egocentric context-aware programming in ad hoc mobile environments. In *Proc. of the 10th Int'l. Symp. on the Foundations of Software Engineering*, November 2002.

[10] N. Minsky, Y. Minsky, and V. Ungureanu. Safe tuplespace-based coordination in multi agent systems. *Journal of Applied Artificial Intelligence*, 15(1), January 2001.

[11] A. L. Murphy, G. P. Picco, and G.-C. Roman. Lime: A middleware for physical and logical mobility. In *Proc. of the 21st Int'l. Conf. on Distributed Computing Systems*, pages 524–533, 2001.

[12] National Center for Supercomputing Applications, Integrated Decision Technologies Group. SAMCat: A securable active metadata catalogue. 2002.