Managing Multiple and Dependable Identities

Digital management of multiple robust identities is a crucial issue in developing the next generation of distributed applications.

Ernesto Damiani, Sabrina De Capitani di Vimercati, and Pierangela Samarati
University of Milan

Our daily activities increasingly rely on remote resources and services — specifically, on interactions between different, remotely located parties. Because these parties might (and sometimes should) know little about each other, *digital identities* — electronic representations of individuals' or organizations' sensitive information — help introduce them to each other and control the amount of information transferred.[1]

In its broadest sense, identity management encompasses definitions and lifecycle management for digital identities and profiles, as well as environments for exchanging and validating such information. Digital identity management — especially support for identity dependability and multiplicity — is crucial for building and maintaining trust relationships in today's globally interconnected society.[2,3]

In this article, we'll investigate the problems inherent in identity management, particularly emphasizing the requirements for multiplicity and dependability. This article's content is a result of our activities performed on the RAPID project (the Roadmap for Advanced Research in Privacy and Identity Management; www.ra-pid.org).

Definitions and Requirements

The term digital identity often refers to two (non-disjoint) concepts: *nyms* and *partial identities*. Nyms give users different identities to use when interacting with other parties in different environments. Behind a nym strong authentication tools such as tokens, smart cards, digital certificates, or biometrics associate individuals with their true digital identities. Weakly bound or unbound nyms, such as those used in peer-to-peer (P2P) file-exchange systems,[4] are meaningful only in the framework of a particular system

 Published by the IEEE Computer Society 1089-7801/03/$17.00

or within a single transaction.

Partial identities are any subset of the properties associated with users (such as name, age, credit-card number, or employment) that the user can select for interacting with other parties.[2] Figure 1 shows an example of partial identities and their use. A partial identity can be named or unnamed, which means it might or might not be related to the user's true identity.

Many of the research issues discussed in this article apply to both partial identities and (strongly bound) nyms, but defining and distinguishing between them is key to describing potential digital identity management solutions. Generally speaking, a digital identity solution should support at least three requirements:

- *Reliability and dependability.* A digital identity must protect users against forgery and related attacks while also guaranteeing to other parties (such as suppliers and brokers in an e-business transaction) that the users can meet transaction-related obligations.
- *Controlled information disclosure.* Users must have control over which identity to use in specific circumstances, as well as over its secondary use and the possible replication of any identity information revealed in a transaction.
- *Mobility support.* The mobile computing infrastructure must be able to take into account its own peculiarities (such as limited bandwidth and display size) to apply multiple and dependable digital identity (MDDI) technology successfully.

In the remainder of the article, we'll examine these requirements in detail to describe a coherent and complete set of functionalities.

Motivations Underlying MDDI Demand

The basic tools underlying MDDI technology have been available for a long time,[5] but the widespread demand for multiple identities is a more recent phenomenon. People conduct a growing percentage of business online, but nearly all of them use physical identities, sometimes guaranteed via digital certificates.[6,7]

Although cost reduction has motivated conventional identity management technology, it plays a smaller role in identity multiplicity and dependability, both of which pose several additional requirements that will likely increase management costs in the short term. However, MDDI technology promises to bridge the existing trust gap between the majority of users and the electronic market, which we hope will boost commercial transactions.

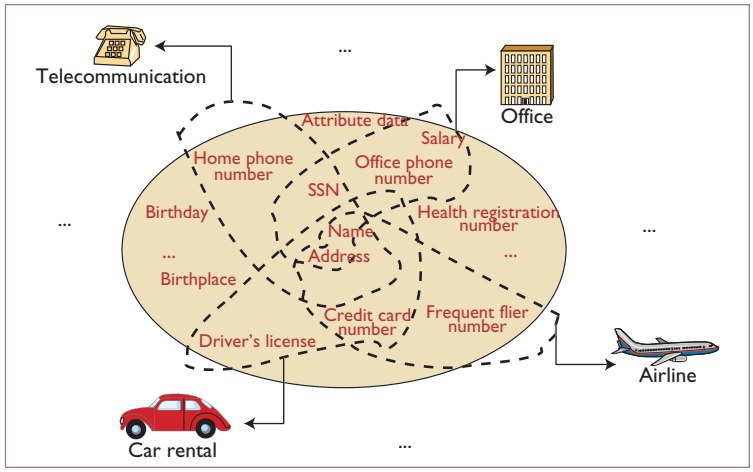

Figure 1. Examples of partial identities. Each dashed line delimits a subset of the user's attributes that can be used as a partial identity when interacting with a party such as an airline or a car rental company.

Two key factors have historically hindered extensive adoption of MDDI solutions: high investment costs (found in PKI-based solutions) and insufficient legislation in the regulatory framework. Today, both factors have developed into powerful driving forces behind most identity management technologies and projects.

PKI-Based Solutions

Until recently, the number of standard digital certificate solutions relying on heavyweight public-key infrastructures (PKIs) has prevented extensive MDDI deployment.[8] The traditional PKI model is rather expensive to set up and manage; when coupled with the low success rate of PKI-based projects, this helps explain industry's reluctance to adopt it. PKI use is still limited either to expensive high-end projects (such as Web-based remote banking or tax collection) or to small pilot efforts (such as those in healthcare). However, some of the more recent managed PKI environments reduce this burden by partially or totally outsourcing digital certificate management — for example, via high-performance appliances (see www.singlesignon.net).

Corporation-wide requirements for cross-platform integration of company services are another motivation toward MDDI. These requirements will likely be less important in the short term, however, because they focus on increased functionality rather than cost reduction.

The common perception in the business community is that MDDI is based on a network of

trusted third parties (TTPs) or other providers of identity authentication and management. This notion suggests a close architectural relationship between PKI and MDDI; for this reason, two recently developed digital identity systems represent part of a new wave that focuses on providing basic functionalities rather than the dependability requirements PKI already satisfies.

Microsoft .NET Passport and Novell DigitalMe. The Microsoft .NET Passport (www.passport.com) infrastructure, together with Novell's less-widespread DigitalMe system (www.digitalme.com), work well as a proof-of-concept pervasive identity management architecture. Although centralized and functionally limited in many respects, they smoothly associate a unique ID with every user and eliminates the need for remembering multiple IDs and passwords for online services. It does this without requiring the specialized infrastructure investment typically found in traditional or managed PKI solutions. This centralized approach also guarantees at least some dependability, which has increased user and enterprise awareness of dependability's importance in e-business and P2P information exchange.

Liberty Alliance. The multinational, multi-industry Liberty Alliance (www.projectliberty.org) consortium is collaboratively developing a set of open standards for federated network identity. The Alliance's specifications build on OASIS's Open Standard Security Assertion Markup Language (SAML, www.oasis-open.org/committees/security), an XML-based security standard that provides a way of exchanging user authentication information. The Alliance chose to extend SAML to include additional security enhancements that are important to identity management, such as opt-in account linking, simple session management, and global log-out capabilities. The Liberty Alliance's openness hints at secure and reliable authentication across hardware and software platforms, including mobile and handheld devices.

Regulatory Framework

The relatively slow adoption of legislation regulating digital identity creation and use has further hindered large-scale MDDI development and deployment. Today, privacy-related legislation is a powerful driver toward adopting digital identities — specifically, support for multiple identities in complex e-business transactions. The US Congress passed the Gramm-Leach-Bliley Act in 1999 to protect privacy data in financial transactions; HIPAA (Health Insurance Portability and Accountability Act) established US regulations about healthcare patient identity privacy in 2000. Following the European Data Protection Directive 95/46/EC, the Directive on Privacy and Electronic Communications 2002/58/EC, and the Electronic Signatures Directive 99/93/EC, European Union member states' legislation has increasingly acknowledged citizen concerns for privacy as well as their reluctance to provide extensive information about themselves during e-business transactions.

Companies around the world increasingly regard MDDI as a viable solution to their difficulties in dealing with legislation regulating consumer data collection and management because MDDI systems can be easily tailored to changing regulations. When dealing with consumer information, companies must consider the impact of both general government policies and those of the countries in which they conduct business. Some regulations even consider the enterprise to be responsible for employees' privacy protection whenever the company interacts with third parties on their behalf.

MDDI System Development Issues

Many concerns about information management have emerged among the actors — namely, individuals (employees, partners, and customers) and e-businesses — that are supposed to benefit from an identity management system. Before we can provide the support needed for multiplicity and dependability in identity management, we must overcome several major issues, as illustrated in Figure 2.

Identity Lifecycle Management

Effective identity management solutions require careful design of the digital identity's full life cycle. Currently, the cycle is modeled as a sequential multiphase process, moving from a creation to a termination phase with support for updating and maintenance. However, such sequential processes do not meet the requirements that multiple dependable identities pose. We can define a novel MDDI-oriented life cycle as a structured asynchronous process that enables co-instantiation and joint evolution of all information items needed to support individuals in different interactions with organizations and to manage the digital identity. MDDI lifecycle management continues to face three open issues.

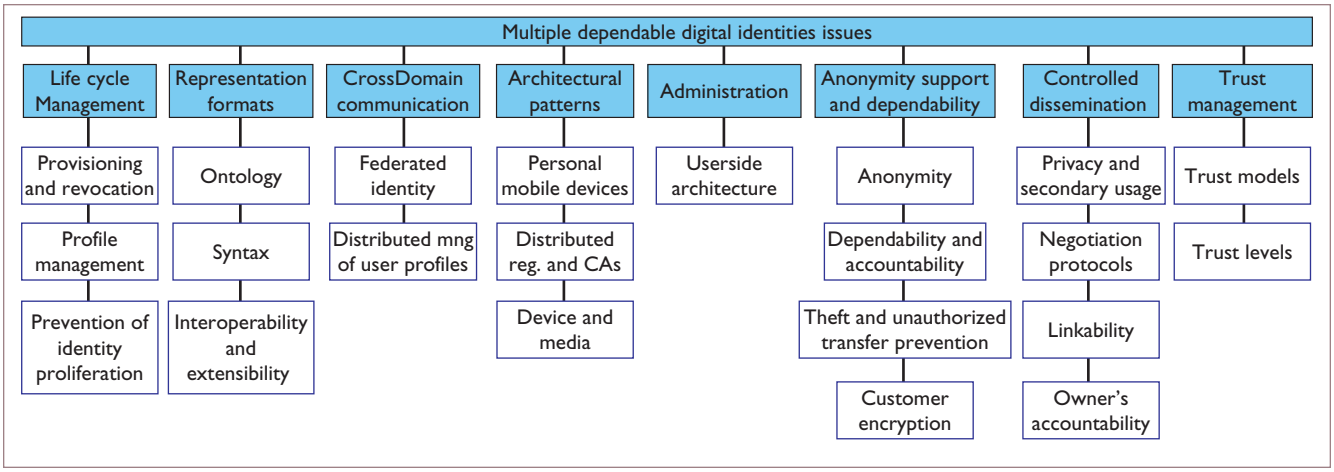

Figure 2. Multiple and dependable digital identity (MDDI) system design issues. The shaded boxes represent the main categories of problems and the clear boxes the specific issues to be addressed.

Provisioning and revocation. Performance is paramount when designing systems that let users efficiently obtain or create identities. One way to achieve better performance would be to automate the process for giving users fast access to information resources. We also need to define ways for automating and securing the process of terminating existing identities, thus preventing improper access on the basis of expired credentials. When employees quit an organization, for example, employers should revoke identity information related to their employment.

Profile management. Users should be able to access and maintain their own identity information and manage their own profiles, but this data's nature and sensitivity might dictate many different solutions. Some data (such as email address) might give the user complete authority; other data might require assistance from third parties to vouch for correctness (such as membership information).

Prevention of identity proliferation. Although some models associate virtually no cost with identity creation, digital identities are indeed resources whose uncontrolled proliferation is highly undesirable. Some researchers have proposed various techniques for transparently imposing soft or hard limits on the number of identities associated with a single individual,[9] but we still lack a comprehensive solution for all situations.

Digital Identity Representation Formats

Another open issue concerns the specification of how identity information should be represented and exchanged. An identity management service must support vocabulary definitions for identity attributes as well as for the control structures used in the protocol itself.

Identity ontology. A fundamental assumption underlying all business transactions is that the involved parties share a common ontology expressing a model of the domain in which the transaction occurs. In the e-business framework, such ontologies are increasingly represented via Semantic Web metadata formats such as Resource Description Format (RDF) and RDF Schema (RDFS).[10] However, we also could use interdomain task ontologies that specify general concepts such as *customer* or *supplier* across multiple domains. Ontology modeling for MDDI and an RDF-like metadata representation for identity assertions are vital for sharing a common concept of digital identity throughout an electronic marketplace. They are also important for establishing relationships between identity and other related concepts, whether general (*person*) or application-specific (*patient*).

Identity syntax. Profile-based digital identities defined from named hierarchical profiles have received increased interest recently. Such identities must satisfy two specifications:

- *Attributes and credentials.* User profiles can be based on existing directory or certificate standards such as the lightweight directory access protocol (LDAP) or SAML. Alternatively, credentials — that is, any user- or third-party-asserted information — can comprise them. Examples are the user's personal data, privileges, habits, or biometrics.

- *Minimal disclosure*. Recent XML-based standards such as SAML provide a uniform syntax for profile assertions. However they use "classical" credentials, such as Kerberos tickets or X.509 certificates,[6] which were designed for authenticating well-known identities. Modern privacy-friendly credentials are promising because they reveal only transaction-specific information.[11,12]

Corporations tend to prefer to build identity structure by collecting user data from a range of available sources, including human resources databases, business workflow, and the users themselves. This identity-collection phase is aimed at automatically or semiautomatically generating user attributes according to business rules.

Identity interoperability and extensibility. Identities must be expressed in a common interchange format, which means the identity management service supports extensible mapping between identities. Moreover, because there is (in principle) no limit on the attributes associated with an identity, identity management must provide for an extensible vocabulary. Again, XML-based data representation formats help here because they can support this extensibility.

Cross-Domain Identity Communication

Enabling reliable communication across multiple user identities in different domains requires several innovative features. Consider a scenario in which a user provides an identity consisting only of a frequent flier number to an online travel agency. The travel agency gives this information to an airline, and then the airline grants the user permission to execute the required action. Regulating who should decide which data composes an identity and assessing how much each partner can trust the assertions provided at each step is an important unsolved problem that requires some key functionalities.

Federated identity management support. In federated identity management systems, a trusted third party supports and manages users' identities. But wide-scale adoption of such systems requires further investigation into techniques for identity-composition and exchange. Even in federated identity management systems, users should be entitled to keep some control over their identities without the burden of protecting and managing them. A particular challenge is the complete retrieval of all needed data while still preserving sensitive information's privacy.

Distributed management of user profiles. As you can guess, the different pieces of information forming an identity can be distributed among several different parties. We still need solutions for retrieving this data and an efficient way to distribute user profile information across different contexts. Modifications to remotely stored identity information must be supported, including the application of update-propagation techniques.

Architectural Patterns for Multiple Identity Management

Today's e-business architectures are based on the old (centralized) PKI concept and must be adapted to the more modern concept of *trust management*. Centralized identity management techniques delegate identity provisioning, maintenance, and revocation to a TTP, which could also be charged with keeping track of the link between users' multiple identities and their single physical identities. Although centralized architectures are robust and simpler to design and implement, they do not fit all MDDI application requirements. Also, once TTPs manage digital identities for large numbers of users, MDDI management systems must scale out to support the voluminous data that large user populations produce. Therefore, we must develop and test new architectural patterns.[13] Hybrid solutions look particularly promising for robustness and efficiency, but we must take many architectural factors into account.

Personal mobile devices. The growing popularity of mobile devices and the current convergence of telecommunication and computer network technologies will produce a broad range of new personal mobile devices that reflect the multiple roles people fill in their daily lives. We need techniques for efficiently and securely storing identity and profile information on such devices.

Distributed registration and certification authorities. MDDI management should not always rely on a single authority; it should flexibly support (hierarchical or P2P) cooperation and interoperability of multiple registration authorities and trust networks.

Devices and media for identity support. MDDI requires dependable devices, such as biometric readers, smart cards, or secure cellular telephones, on which users can divulge their identities. Users should be able to choose different methods depending on the application.

Identity Administration

Maintaining multiple identities as separate and independent named sets of attributes or credentials obviously poses huge management problems. Therefore, MDDI solutions should provide protocols, tools, and techniques for fast and reliable credential updates. Like database views, profile-based digital identities can be materialized (actually stored) or virtual (computed at runtime: only the definition is stored) according to application needs. An open research issue is how to achieve seamless and scalable view-computation over profiles by selectively using partial encryption or profile data transformation. Alternative approaches rely on hidden attributes that users can disclose selectively via multiple encryption keys. Such disclosure lends itself to user-side execution, whereas dynamic view-computation can be executed only on a server. Identity management solutions should also be integrated with personalization solutions to allow profile reuse. Ideally, the development of a user-side architecture should also help the user decide the consequences of releasing certain information.

Anonymity Support and Dependability

MDDI must provide users the ability to remain anonymous if identity information is not required in a transaction. Anonymity does not imply that no information be released, but it requires that any released information be unidentifiable. The challenge is that the information, although anonymous, must be proven reliable.

Anonymity support. An anonymous communication infrastructure is a basic requirement for any secure, privacy-friendly identity management system; if identifying Internet users simply by looking at their IP addresses is easy, then any effort to provide privacy-protecting mechanisms at a higher level is pointless. Nevertheless, some degree of anonymity support — possibly coupled with reputation and trust provisions — is essential for future MDDI solutions, which means we must address several important topics. Ideally, an MDDI service should support different degrees of privacy in transactions. In particular, parties should be able to remain completely anonymous or declare some identity information about themselves, including weakly bound pseudonyms, which may or may not provide linkability. (From a legal perspective, the term *anonymity* should be reserved for situations in which no linkability is possible, but to be consistent with existing technical literature, we use here a loose definition of the concept in which we can consider different degrees of identity protection.)

Dependability and accountability. Dependability and accountability techniques are needed for guaranteeing every party's ability to meet its obligations in an e-business transaction. One way to achieve this goal while preserving anonymity is to integrate multiple identities with secure audit trails, which can keep track of accessed resources. In distributed business processes, audit information must be shared among all interested parties to hold them accountable. Moreover, it must be possible to follow an audit trail in cases of security breaches and identity misuse. Identity management systems should enforce good audit practices and support forensic analysis.

Theft and unauthorized-transfer prevention. Identity fraud affects users negatively at both home and work. For this reason, early models for digital identities addressed theft prevention and, in general, described countermeasures to improper digital identity use. However, several problems remain to be solved related to how an identity management system should guard against identity theft via automating provisioning, maintaining, and revoking digital identities.

Custom encryption techniques. MDDI management must be based on underlying cryptographic techniques that provide a degree of trust about identity correctness and protect against misuses. Possible encryption techniques tailored to anonymity support include

- *Pseudonym systems*. Each digital identity is associated with a pseudonym, and any pseudonyms that refer to the same individual should remain unlinkable from each other. Although current privacy-enhancing technologies can prevent linkability, such an ability is insufficient for creating a privacy-friendly trading environment that supports multiple identities. For instance, the information shared among the different phases of a single e-business transaction could be used to circumvent unlinkability, and yield a detailed user profile.
- *Cryptographic credentials*. Traditional credentials are ill suited for MDDI environments because they always provide the same amount of information, regardless of the specific transaction. Modern privacy-friendly credentials are more promising for MDDI because they reveal

only the information that is strictly necessary to perform a transaction.
- *Group signature schemes.* Standard digital signatures are linked to physical identities via digital certificates. Group signature schemes support weak anonymity to signers by hiding them behind group memberships.

Many other cryptographic solutions exist today that are potentially suitable for MDDI.[14] Researchers still need to evaluate their applicability and performances in MDDI-specific scenarios and determine how to use them in combination with MDDI identity management tools. There is also a need for fundamental research to ensure that alternative identity-management approaches are available in the event of unexpected successes in breaking factoring-based cryptography or in developing quantum computing.

Controlled Dissemination of Authenticated Information

As described earlier, MDDI technology must operate in an environment in which well-defined trust models underlie the interactions among all parties involved. Users must have control over which identity to use and which attributes to disclose to their counterparts.

Privacy and secondary usage control. Although identity attributes should be enriched with privacy preferences, languages for expressing privacy requirements as part of an identity are still in their infancy. As a first step in that direction, there are technologies, such as A P3P Preference Exchange Language (APPEL),[15] that support the platform for privacy preferences (P3P). We also need approaches that let parties evaluate their counterparts' privacy policies for secondary use. The identity management service should be able to support fine-grained restrictions on the release of identity information and, possibly, to associate access and usage restrictions that must be obeyed with the released information.

Negotiation protocols. To maintain control over the release of identity information, researchers must develop approaches for preventing cases in which an identity is released but no service is given in return. We still need negotiation protocols that let parties determine which identity information to release to their counterparts.[1]

Linkability. Although users must be able to choose the identity they wish to adopt when interacting with other parties, researchers cannot leave entirely to the user agent the choice of which attributes or credentials to associate with specific identities. A user can employ different identities or pseudonyms in different transactions, but these transactions are still linkable, thus allowing user profiling, which can put true identity at risk. We need innovative methods for controlling possible linkage between different identities. Group sharing solutions can help to diminish linkability threats.

Owner's accountability. Identity owners might need to trust automated agents to correctly represent their identities during e-business transactions. Identity management services should provide adequate accountability to let users track how their identity information is managed and to whom and in what context it is released.

Trust Management

In a federated identity management system, users can get credentials in many different ways, but entities trying to verify such credentials often have no direct means of assessing their trustworthiness. From the user's viewpoint, one key problem with today's Web systems is the requirement for a different password on each system. Password management addresses this problem; current solutions let users access multiple distributed resources with a single sign-on facility, but the systems rely on other partners to trust the authentication process used to approve the identity credentials.

Trust management in a privacy-enhancing MDDI environment means defining methods for receiving reliable evidence about credentials and assessing their degree of trustworthiness. In principle, strong authentication techniques such as tokens, smart cards, digital certificates, or biometrics can guarantee trust in a digital identity. The current trend is toward providing Internet-based trust services, which deal with various aspects of trust and are held accountable for the services they provide. Another aspect of confidence and trust is related to the ability to evaluate and assess the security level of the components, systems, and services used to authenticate a user or to relay authentication information.

Trust models. We need innovative models that identify the conditions under which a party can trust others for security and privacy.[16] For instance, developers can devise reputation models that allow reliable associations between reputa-

tions and nyms.[17] To develop reputation models that allow trust establishment in anonymous or pseudoanonymous systems, researchers must develop techniques and protocols for measuring reputations, storing and sharing reputation information, and providing reliability.

Support of trust levels. How much should a party involved in a business transaction trust another party's digital identity? Different levels of trust should be possible. For example, users could directly provide nonsensitive information (low trust), or they could provide certificates that verify their identities (trust levels would depend on the certificate authority's credibility).

Conclusions

Some of the issues we covered here are likely to have much greater impact than others. Providing a tunable degree of anonymity via flexible identity management, for example, could lead to new application areas in P2P and grid computing. Traditionally, "pure" P2P approaches to MDDI were aimed at creating, managing, and exchanging digital identity information without intermediaries, but this nearly always produced some kind of weak anonymity that proved wholly unfit for business applications. The next generation of MDDI systems will probably solve this problem by returning full control to users over the strength of their nym's binding. As a precursor of this vision, Ping ID (www.pingid.com) gives all users the option between exchanging (weakly bound) nyms directly through P2P protocols and using Web technology to authorize a trusted third party to dispense (strongly-bound) identity data.

Sound modeling of this flexibility and incorporating it into core P2P development architectures such as Sun's JXTA (www.jxta.org) are challenges for both academic and industrial research. Future personal (and multiple) digital identities will be stored in several places, from the user's workstation to remote service-provider-managed repositories. Achieving the research goals outlined in this article will enable a new generation of advanced MDDI services on the global information infrastructure.

Acknowledgments

The work reported in this paper has been partially supported by the EU-funded Roadmap for Advanced Research in Privacy and Identity Management (RAPID) and by the Italian Ministry for Research under the KIWI and MAPS projects.

References

1. P. Bonatti and P. Samarati, "A Unified Framework for Regulating Service Access and Information Release on the Web," *J. Computer Security*, vol. 10, no. 3, 2003, pp. 241–272.
2. S. Claub and M. Kohntopp, "Identity Management and Its Support of Multilateral Security," *Computer Networks*, vol. 37, 2001, pp. 205–219.
3. B. Pfitzmann, "Privacy in Enterprise Identity Federation: Policies for Liberty Single Signon," *Proc. Workshop on Privacy Enhancing Technologies*, Springer Verlag, 2003.
4. A. Oram, ed., *Peer-to-Peer: Harnessing the Power of Disruptive Technologies*, O'Reilly & Associates, 2001.
5. D. Chaum, "Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms," *Comm. ACM*, vol. 24, no. 2, 1981, pp. 84–88.
6. R. Housley et al., "Internet x.509 Public Key Infrastructure Certificate and CRL Profile," RFC 2459, Internet Eng. Task Force, Jan. 1999; www.ietf.org/rfc/rfc2459.txt.
7. J.S. Park and R. Sandhu, "Binding Identities and Attributes Using Digitally Signed Certificates," *Proc. 2000 Ann. Computer Security Applications Conf.*, IEEE Press, 2000, pp. 120–127.
8. S. Brands, *Rethinking Public Key Infrastructures and Digital Certificates: Building in Privacy*, MIT Press, 2000.
9. A. Buldas, P. Laud, and H. Lipmaa, "Accountable Certificate Management Using Undeniable Attestations," *Proc. 7th ACM Conf. Computer and Communications Security*, ACM Press, 2000, pp. 9–17.
10. "Resource Description Framework (RDF) Model and Syntax Specification," W3C note, Feb. 1999; www.w3.org/TR/REC-rdf-syntax.
11. J. Camenisch and E. Van Herreweghen, "Design and Implementation of the Idemix Anonymous Credential System," *Proc. 9th ACM Conf. Computer and Communications Security*, ACM Press, 2002, pp. 21–30.
12. T. Yu and M. Winslett, "Policy Migration for Sensitive Credentials in Trust Negotiation," *Proc. Workshop on Privacy in the Electronic Soc.*, ACM Press, 2003, pp. 9–20.
13. E. Damiani, S. De Capitani di Vimercati, and P. Samarati, "Towards Securing XML Web Services," *Proc. 2002 ACM Workshop on XML Security*, ACM Press, 2002, pp. 90–96.
14. D. Boneh and M. Franklin, "Identity-Based Encryption from the Weil Pairing," *SIAM J. Computing*, vol. 32, no. 3, 2003, pp. 586–615.
15. L.F. Cranor, *Web Privacy with P3P*, O Reilly & Associates, 2002.
16. M.N. Huhns and D.A. Buell, "Trusted Autonomy," *IEEE Internet Computing*, vol. 6, no. 3, 2002, pp. 92–95.
17. S. De Capitani di Vimercati et al., "Managing and Sharing Servents Reputations in P2P Systems," *IEEE Trans. Knowledge and Data Eng.*, vol. 15, no. 4, 2003, pp. 840–854.

Ernesto Damiani is a professor in the Department of Informa-

tion Technology at the University of Milan and serves as the vice chair of the ACM Special Interest Group on Applied Computing (SIGAPP). His research interests include distributed and object-oriented systems, semistructured information processing, and soft computing. He has a PhD in computer science from the University of Milan. Contact him at damiani@dti.unimit; www.dti.unimi.it/~damiani.

Sabrina De Capitani di Vimercati is an associate professor in the Department of Information Technology at the University of Milan. Her research interests are in the area of information security, databases, and information systems. She received a PhD in computer science from the University of Milan. Contact her at decapita@dti.unimi.it; www.dti.unimi.it/~decapita.

Pierangela Samarati is a professor in the Department of Information Technology at the University of Milan. Her research interests are in data and application security and privacy, access control policies, models and systems, and information protection. She is the chair of the IFIP (International Federation for Information Processing) working group on data and application security (WG11.3) and a member of the steering committee of the ACM Special Interest Group on Security, Audit, and Control (SIGSAC). Contact her at samarati@dti.unimi.it; www.dti.unimi.it/~samarati.