Context-Aware Middleware for Anytime, Anywhere Social Networks

Dario Bottazzi, Rebecca Montanari, and Alessandra Toninelli, *University of Bologna*

The Socially Aware and Mobile Architecture lets you create anytime, anywhere social networks among users in physical proximity. Samoa exploits semantic-based context modeling and matching algorithms for social-network extraction.

Sociality characterizes an individual's life: people go to bars and restaurants, study together in schools, and work in teams on production lines and in business. Every person plays a role within a social community. Social ties, such as friendship, common interests, and shared professional activities, bind individuals together. This web of social bindings is referred to as a *social network*.[1]

Since the 1960s, research and industrial efforts have investigated advanced collaborative systems for leveraging human connections and improving human interactions in workspace environments. Although the computer-supported cooperative work field has contributed much in this direction, CSCW solutions generally focus on interactions driven by business needs, where connections among people tend to be formal and structured. Only recently has the convergence of social and computing disciplines focused attention on the design of *social-networking services*—applications that support human social interactions and are characterized by their swarming, transitory, and informal qualities.[2–4]

Technology advances in wireless networks and the increasing diffusion of portable devices offer a unique chance to improve social-networking services. The formation of ad hoc networks enables serendipitous social encounters between proximate users with common interests, anywhere and anytime.[2,5,6] These services depend on ubiquitous technologies to shift the application focus from virtual to physical social spaces. Physical proximity increases the likelihood of impromptu social relationships. Physical places can also act as social filters. Museums and discos, for example, group together people who are likely to share common characteristics and preferences. Several recent prototypes exploit individuals' colocation and reciprocal proximity for guiding social-network formation and management strategies and for restricting the scope of user interactions.[5,6]

However, various technical challenges remain. Anytime, anywhere social computing requires several support mechanisms and tools, including location and proximity systems, expressive representation models of physical place and user characteristics, and effective social-matching algorithms. Moreover, because of the impromptu and transient nature of ubiquitous interactions, these solutions must minimize user intervention. Anytime, anywhere social-network computing also requires shared and interoperable vocabularies for modeling location and entity characteristics. Current solutions tend to address only a subset of these issues.

We believe that the success of anytime, anywhere social computing depends on middleware solutions that separate social-network management concerns from application requirements. Our middleware solution, the Socially Aware and Mobile Architecture (Samoa), integrates a set of common management facilities for personalizing location-dependent social networks, and for propagating social networks' visibility up to the application level.

1541-1672/07/$25.00 © 2007 IEEE
Published by the IEEE Computer Society

Emerging requirements and solution guidelines

Several Internet-based social-networking prototypes address various application domains. These include the automatic identification of acknowledged experts in a field,[7] the study of social dynamics characterizing a certain social system,[8] and the improvement of search engine ranking algorithms.[9] These applications share an underlying assumption that the Internet is the main source of empirical data from which to extract social relations and patterns. Web pages,[9] scientific publication databases,[10] XML-based friend-of-a-friend files,[11] personal or organization mail repositories,[8] and Usenet posts[7] are examples of commonly used information sources. Most current Internet-based social systems adopt social-network analysis approaches that are based on data mining techniques:[7,12] co-occurrence of names in a Web page,[9] citation or coauthoring of a scientific article,[10] and sender and recipient of an email.[8] Only recently have researchers proposed ontology-based approaches for modeling and inferring relevant social relationships among individuals.[11]

Social-network systems for ubiquitous-computing environments are an emerging trend in social computing. However, because ubiquitous-computing environments are more dynamic and heterogeneous than Internet-based deployment scenarios, they require solutions that follow different design guidelines from their Internet-based counterparts. The set of potentially available social-network members varies because of user mobility, so it can't be predetermined. Interactions among individuals are typically opportunistic, transient, and serendipitous, and often involve users in physical proximity. In addition, the large scale and dynamicity of ubiquitous environments make social-network management challenging. It's therefore crucial to identify appropriate criteria for delimiting the searching space for social-network members (the *discovery scope*). Social interests and affinities are too coarse grained, while the physical place where users are likely to establish interactions is an additional first-level criterion for social-network extraction.[5,6]

The need to shift to a location-centric social-computing paradigm has guided the design and development of recent social-networking solutions. Systems like LoveGety, ProxyLady, and SocialNet exploit proximity or colocation visibility to verify affinities among colocated users.[6,13]

However, current solutions are still more proof-of-concept application prototypes of single management aspects than comprehensive frameworks for supporting the design, development, and deployment of anytime, anywhere social-networking services. To our knowledge, all literature proposals are built on top of the network layer and tend to provide dedicated support for specific applications. This approach has some limitations. First, you can rarely reuse ad hoc support in different application domains, so you must build a new support system from scratch whenever you develop a new application. In addition, building social-networking applications on top of the network layer can be tedious and error-prone because you must deal explicitly with all issues related to user and device mobility, intermittent connectivity, and availability.

> Because ubiquitous computing environments are more dynamic than Internet-based deployment scenarios, they require solutions that follow different design guidelines.

So, we need middleware solutions that address social-network management details, such as user location detection and tracking, user profiling, and social matchmaking. Such solutions would let application developers focus on designing and developing the application logic. This significantly simplifies and accelerates application development. Designers could use the same middleware-level support in different social-computing applications, thus encouraging applications' interoperation and rapid prototyping.

To support the creation of social networks that reflect the reality of social interactions in ubiquitous environments, we must account for context information, such as user location and reciprocal proximity, user attributes, motivations, attitudes, activities, and social preferences.[12] Toward this goal, middleware proposals should provide integrated support for context modeling, acquisition, reasoning, and context-aware social-network extraction. In particular, we need adequate expressive means for rich and unambiguous representation of users, their contexts, and the networks they participate in.[12] However, the impossibility of making a priori assumptions about how user contexts are described in an open and dynamic deployment scenario complicates context-modeling endeavors. Semantic Web languages seem to offer a promising solution to the issue of describing social contexts at the proper abstraction level, while enabling automated reasoning on context representations. In addition, emerging ontology standards, such as the Resource Description Framework (RDF) and Web Ontology Language (OWL), allow interoperability between possibly unknown users who might wish to establish a social interaction.

The Samoa framework

The Samoa framework supports the creation of anytime, anywhere semantic context-aware social networks—that is, the logical abstractions that group together mobile users who are in physical proximity and share common affinities, attitudes, and social interests. In particular, Samoa lets mobile users create roaming social networks that, following user movements, reflect at each instant all nearby encounters of interest. Samoa roaming social networks center on a user (the *ego user*), and are based on two kinds of context visibility:

- *place visibility* (place awareness)—the visibility of the user's physical place, and
- *profile visibility* (profile awareness)—the visibility of place or user requirements and characteristics.

Place visibility restricts the discovery scope for social-network extraction to entities in the same place as the ego user. The visibility of user or place profiles further refines the discovery scope to create personalized social networks. In addition, Samoa models and represents context data in terms of semantic metadata (profiles) and exploits semantic-matching algorithms for analyzing profiles and inferring potential semantic compatibility.

Social-network management model

The Samoa social-network management model defines three management roles:

- *Managers* are the mobile ego users interested in creating social networks. They're responsible for defining the discovery scope boundaries of their social network and the criteria guiding its extraction.

- *Clients* are users located within the discovery scope boundaries and are eligible to become members of the manager's social network.
- *Members* are users affiliated with a social network.

Each mobile user can play all roles. The manager role can be covered by a human or by a software component acting on the ego user's behalf.

In Samoa, social-network management is based on the concept of place, which lets us establish well-defined discovery scope boundaries. As figure 1 shows, each manager defines its own place. The manager is the center of the place, and the place is the set of all Samoa clients who are physically proximate to the manager—that is, those devices that are connected to the manager device by a routing path of a maximum length of *h* network hops, called the *place radius*. In Samoa, network hops represent the distance between two physically connected entities. For example, two entities whose devices are within each other's communication range have a distance of one hop. We don't determine the set of clients in a place a priori; rather, it dynamically changes as users move and devices are disconnected and reconnected.

Depending on the application deployment scenario, different mappings of the place abstraction are possible, either fixed or mobile. For example, a place might define the set of users whose devices are currently connected to the same wireless cell or to the same mobile ad hoc network (Manet). Places can overlap, or they can be defined by more than one manager—for example, two managers could be allocated at a one-hop distance. Users can freely roam among places and might be clients of more than one place at any time.

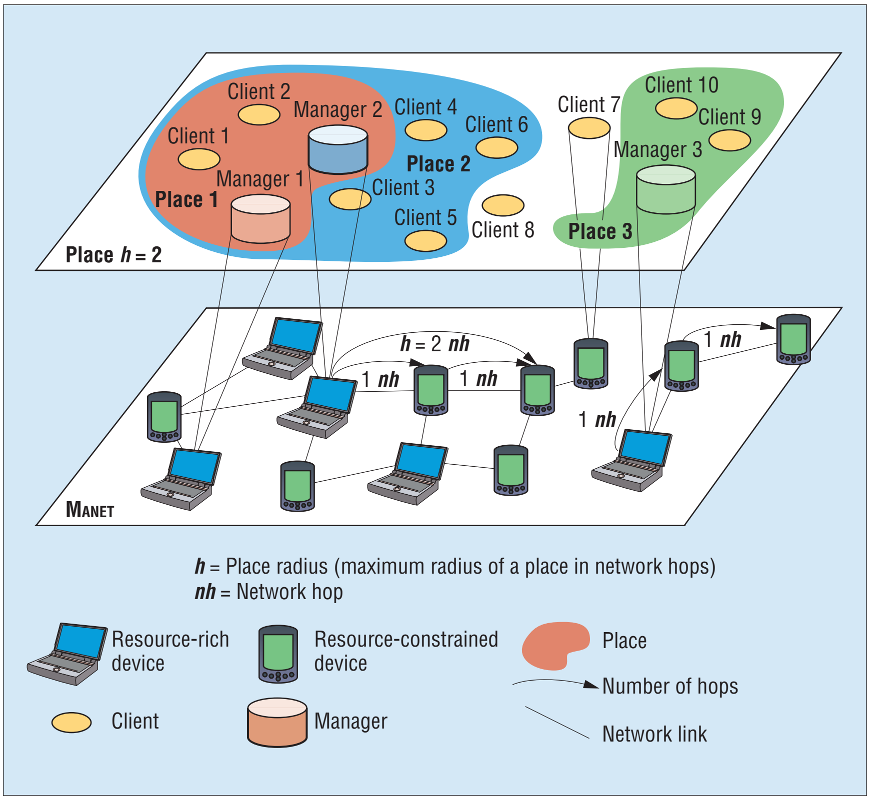

Figure 1. An example place mapping of the socially aware and mobile architecture (Samoa) onto a mobile ad hoc network.

Metadata model

All Samoa entities—that is, places and users—are associated with unique identifiers and profiles describing their characteristics. Profiles have a modular structure comprising different parts, each grouping metadata with a common logical meaning.

A *place profile* has two parts:

- The *identification* part includes a unique identifier, a name, and a description of the physical place.
- The *activity* part includes all of the social activities that characterize the place and that members of that place can share. For example, a bookshop's profile might include activities such as shopping and reading.

Samoa's place profiles account not only for users' social preferences, but also for the relationships between people and the places where they're located and where interactions are likely to occur. The underlying assumption is that the places where users move and operate will influence their activities and interactions with other users.

The *user profiles* for all Samoa users (managers, clients, and members) include identification and preference parts. The identification part provides user naming information, such as a personal identifier, and describes user properties, such as age, gender, and education. The preference part defines user activities—in particular, the activities the user is interested in and, for each of these activities, the user's specific preferences. For example, in the user profile in figure 2a, the user is interested in the shopping activity, and mostly prefers books about history that cost less than 80 euros and are at the superstore Harrod's.

Samoa managers also have a discovery profile associated with each place they manage. The *discovery profile* defines the preferences clients must match to join the manager's social network. Similarly to the user profile preferences, discovery profile preferences include desired client attributes for each activity. For instance, a manager's discovery profile might state that he or she is looking for other users of the same age who are interested in the shopping activity, preferably in buying books.

Social-network extraction model

Samoa lets managers exploit two social networks.

- The *place-dependent* social network shows only the members currently colocated with the network manager.
- The *global* social network persistently records the whole set of place-dependent social networks dynamically created over time as the manager moves across places.

Samoa determines place-dependent social networks through two semantic-matching

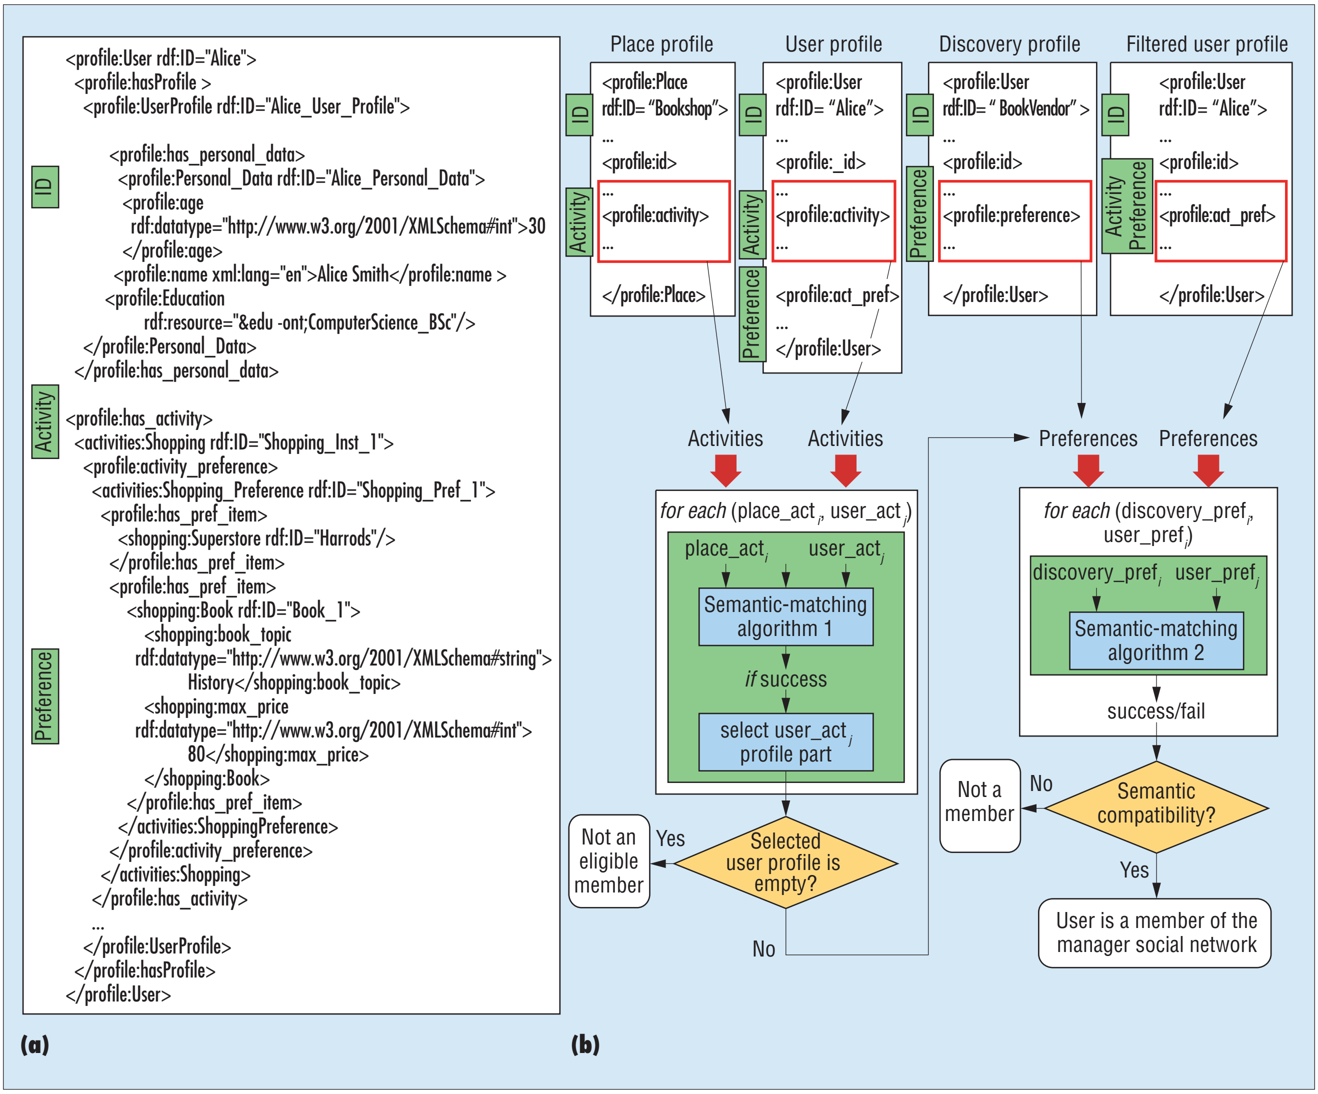

Figure 2. User profiles in Samoa: (a) user profile example and (b) profile-based social-network extraction.

algorithms. The first algorithm operates on user and place profiles to identify a first set of eligible members within a place's discovery scope. It compares all user profiles currently in the manager's place with the place profile. Only those users whose profiles have activities that are semantically related to the place profile activities become eligible members. The algorithm returns the semantically compatible user profile parts of eligible members, as figure 2b shows.

The second matching algorithm selects as members only those users whose attributes semantically match the preferences included in the place manager's discovery profile. In particular, the algorithm iteratively analyzes all user profile parts returned by the first algorithm to determine whether the preferences in the eligible members' user profiles semantically match the preferences in the manager discovery profile (see figure 2b). By applying the matching algorithm to all eligible members colocated in a place, Samoa builds the manager's social network for that place.

As figure 3 shows, both matching algorithms exploit description-logic-based subsumption reasoning to determine whether a particular individual is an instance of a certain class. Toward this goal, we represent place activities and preferences in the manager's discovery profile as classes, and define user activities and preferences in eligible members' user profiles as instances. In addition, we define activity and preference classes by constraining their specific properties to assume a certain value or range of values—for example, we could define a preference class about shopping in the manager's discovery profile by constraining the property representing the purchased object to assume the value "book."

Suppose that the property representing the purchased object in a user preference instance is set to the "book about history" value. In this case, the matching algorithm infers that the user preference is an instance of the preference class in the manager's discovery profile.

The matching algorithms can recognize different semantic relationships on the basis of the subclass relationships defined in the activity or preference ontologies. In particular, the algorithms recognize three semantic similarity relationships. The user activity or preference can be

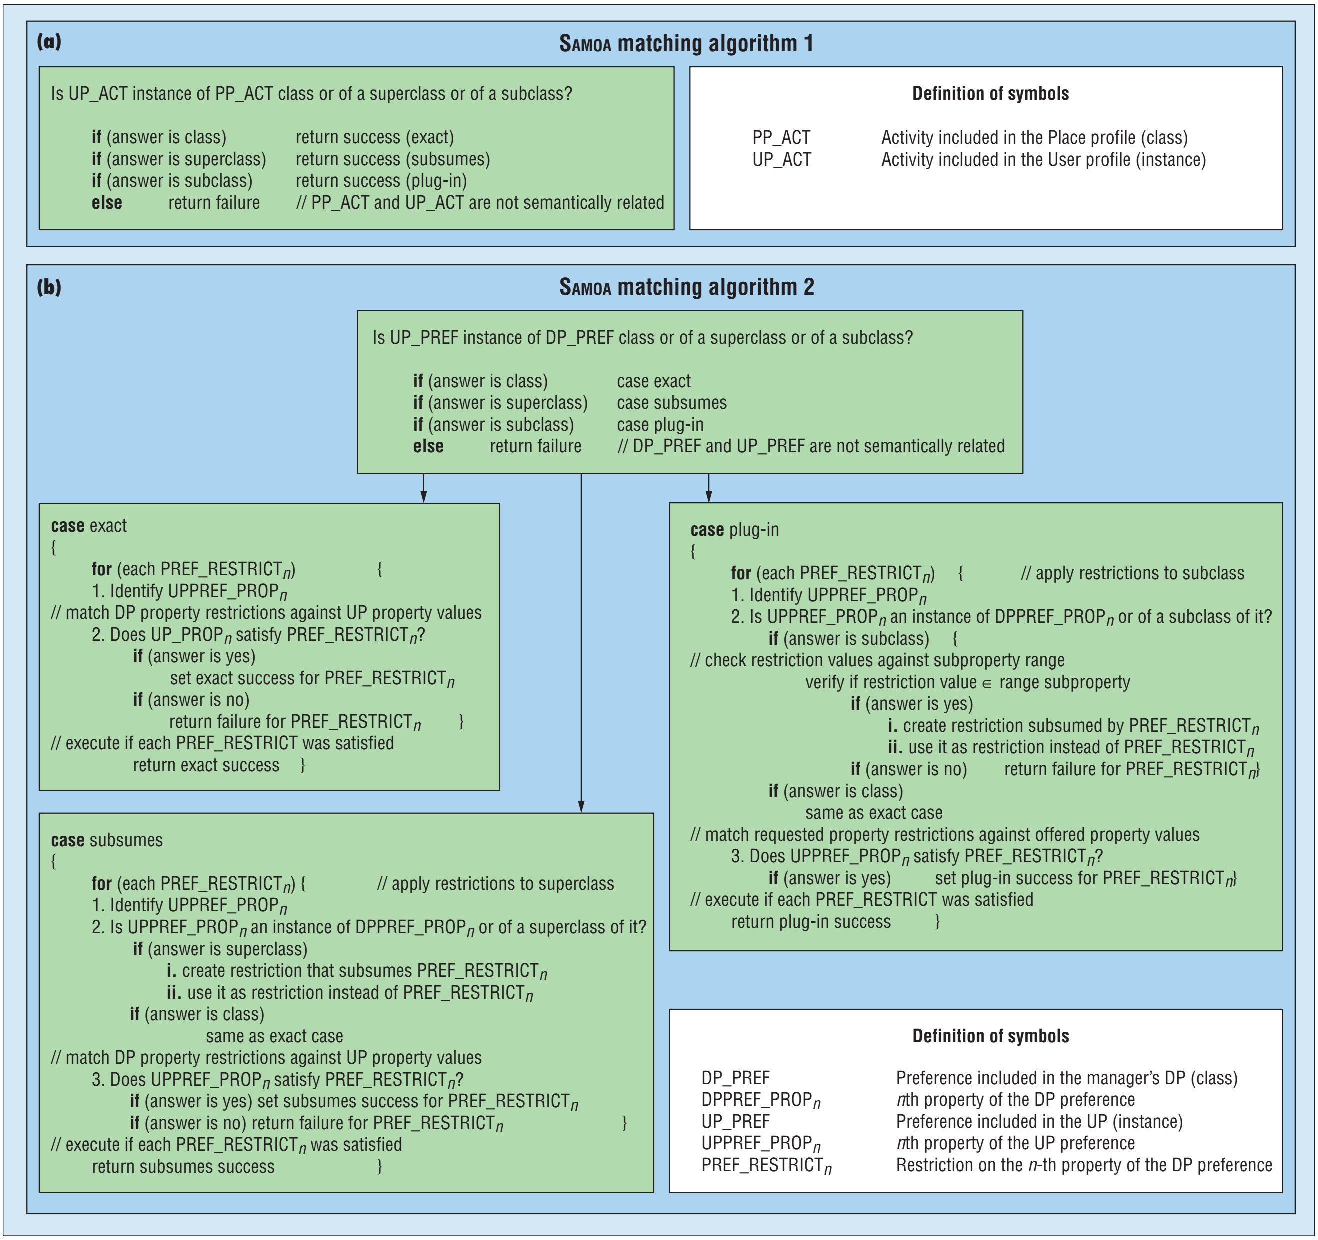

Figure 3. SAMOA matching algorithms. (a) The first algorithm identifies a first set of eligible members in a manager's place. (b) The second algorithm analyzes the first algorithm's results to determine which members' preferences match those of the manager.

- an instance of the activity or preference class in the manager's place or discovery profile (*exact* case), or
- an instance of a more generic activity or preference class (*subsumes* case), or
- an instance of a more specialized activity or preference class (*plug-in* case).

SAMOA incrementally builds a manager's global social network by maintaining information about the members of all transient place-dependent social networks. In particular, we permanently include the profile of any new member in a place-dependent social network in the manager's global social network. Place-dependent social networks let managers easily discover colocated users of interest when they want to establish one-shot and transient interactions. Global social networks, on the other hand, let managers create application-dependent past interaction histories that can enable more complex collaboration strategies and patterns.

SAMOA middleware

SAMOA middleware has a layered architecture. It's built on top of the Java Virtual Machine and is organized in two logical layers: the basic service layer and the social-network management layer.

The basic service layer provides facilities to support naming, detection of colocated SAMOA entities, and communication.

The *message transport manager* implements UDP-based point-to-point and multipoint communication patterns. The MTM point-to-point communication support lets SAMOA entities send messages to a host identified by a known IP. The MTM multipoint communication support lets SAMOA entities broadcast a message to other Samoa entities allocated in the same place by following a broadcast protocol for cell-based environments and a flooding protocol for MANET settings.

The *location/proximity manager* generates and assigns user personal and place IDs by exploiting a naming approach that statistically ensures identifier uniqueness. The L/PM also lets SAMOA entities advertise their online availability by broadcasting advertisement messages at regular times. Advertisement messages include the entity's personal and place IDs and IP address. All SAMOA entities rely on the L/PM to sense incoming advertisements and to build a table of sensed colocated users. Table entries are associated with time stamps. If the L/PM doesn't receive beacons from an entity within a defined threshold, it removes the associated entry and considers the entity device disconnected. The L/PM disseminates advertisements only within the place's scope by coordinating with the MTM. It determines the physical area delimiting a place only once, at place deployment time.

The social-network management layer includes facilities for semantic-based social-network extraction and management.

The *profile manager* (PM) provides graphic tools for specifying and checking profiles for correctness and for distributing user, discovery, and place profiles to interested SAMOA entities and support facilities. As figure 2a shows, SAMOA adopts OWL-based formats for profile representation. The Java-based ontology editor Protégé enables profile visualization and browsing.

The *semantic-matching engine* supports semantic matchmaking according to the two matching algorithms. In the current implementation, the SME relies on the Pellet reasoner's (http://pellet.owldl.com) reasoning capabilities, while it stores and accesses OWL ontologies via the Jena Semantic Web Framework (http://jena.sourceforge.net). The SME accesses the Pellet reasoner via Jena APIs and SPARQL queries.

The *place-dependent social-network manager* manages the place-dependent social network—that is, it creates and maintains a table that includes all manager social-network members that are currently colocated. The PSNM coordinates with the PM to obtain members' personal IDs and user profiles. When a member connects or disconnects from the place, the PSNM updates the table to reflect the change (via coordination with the L/PM). Each PSNM entry includes information about members, such as their personal and device IDs, along with the returned user profile semantically compatible with the place profile.

The *global social-network manager* (GSNM) creates the global social network by maintaining and storing all place-dependent social networks in a dedicated table. Each table entry stores all personal IDs and user profiles of all members in the manager's social network. In addition, the table stores the place and manager discovery profiles that guided each member's selection.

> In a viral-marketing scenario, information forwarding follows a word-of-mouth model based on the customer's encounters during shopping activities.

Viral-marketing scenario

We evaluated the SAMOA framework's performance by deploying it in a viral-marketing scenario. SAMOA support facilities let vendors and customers build social networks to distribute specific product advertisements. Vendors can forward promotional messages (for example, book discounts) to customers currently in their bookshop (using the vendor's place-dependent social network) and to all customers who previously visited their bookshop (using the vendor's global social network). A customer who receives a promotion can use SAMOA to forward the information to nearby customers (using the customer's place-dependent social network) and to all customers previously encountered in all visited bookshops (using the customer's global social network). Information forwarding follows a word-of-mouth model based on the customer's encounters during shopping activities.

Deployment setting

As a testbed scenario for our viral-marketing application prototype, we consider the case of a shopping mall with various shops, including one bookshop. IEEE 802.11-compliant access points provide connectivity. In particular, each shop has one access point that provides wireless connectivity to all customers within the access point coverage area. Customer laptops are equipped with IEEE 802.11b/g wireless cards running the viral-marketing application. A bookshop server hosts the prototype. The server host and customer devices run all SAMOA support facilities.

In our deployment scenario, the book vendor and the customers all play the manager role in that they define their places and create their own social networks. The vendor-defined place is fixed, with the server acting as its center, while the customer's place is mobile. Its allocation is determined by the network cell of the shop the customer is currently attached to—for example, the bookshop cell. Given this deployment setting, the discovery scope for social-network extraction is restricted to SAMOA entities whose devices are connected to the same wireless cell. Vendors and customers appear as clients when in places managed by others.

Programming the application

The realization of the SAMOA-based viral-marketing application has two distinct phases.

- In the *application programming* phase, application developers define and code only the application functionalities, without dealing with social-network management issues.
- In the *application deployment* phase, developers specify application-specific configuration parameters and those SAMOA profiles needed to guide appropriate social-network extraction transparently to the application. At deployment time, the viral-marketing application lets the vendor describe the commercial promotions for potentially interested people, the bookshop place profile, the discovery profile, and the vendor's user profile.

Our prototype promotion descriptions include the shop contact information (name, address, telephone number, and e-mail) and the set of books that are sold at a discount and their prices. Figures 2 and 4 depict some profile examples valid for the application prototype. These include a customer's user profile; the bookshop's place

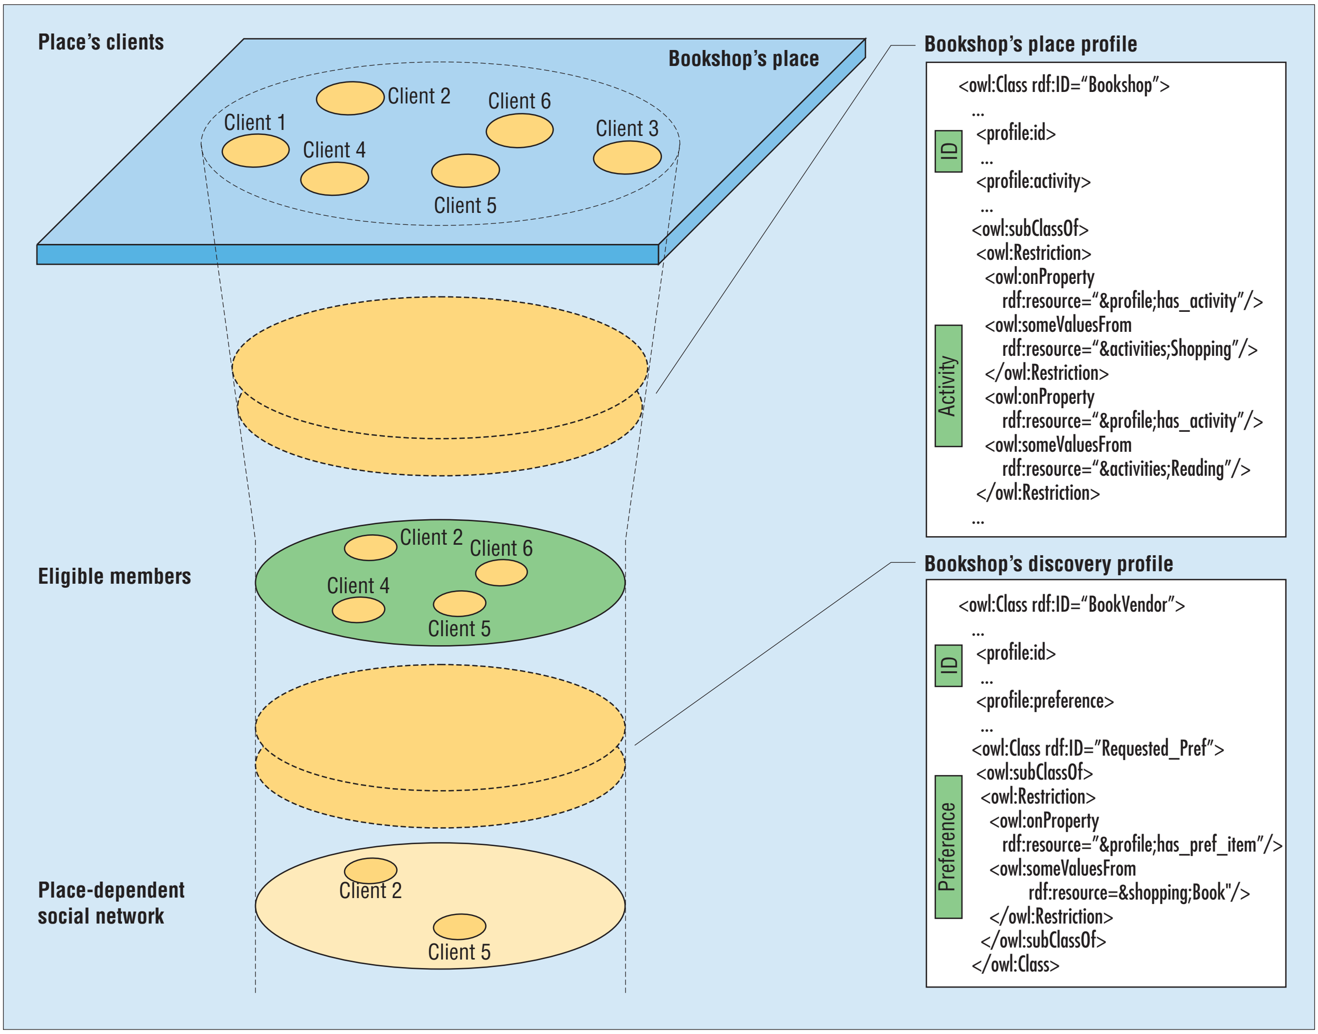

Figure 4. Bookshop's place and discovery profiles and their use in social-network extraction.

profile, describing activities that characterize the bookshop (such as shopping and reading); and the vendor's discovery profile, defining the manager's preference for customers interested only in buying books. Similar considerations apply to the viral-marketing applications running on the customer's devices.

Social-network extraction

To determine their place-dependent social networks, vendors must first identify the place. In our testbed setting, the server acts as the manager and the server's managed place is mapped onto the bookshop wireless cell (see figure 5a). Figure 5b shows how Samoa support facilities interact to build the vendor's place-dependent social network.

When a customer enters the bookshop, the customer's device connects to the locally available wireless cell. The L/PM instance running on the customer's device advertises the customer's availability and detects a new place's availability—that is, the vendor's place. Then, the PM instance running on the customer's device coordinates with the PM component running on the vendor's device to obtain the place profile. When the customer PM receives the profile, it coordinates with the SME facility to filter the customer's user profile according to the bookshop's place profile. If Samoa finds semantic matches, the customer is considered an eligible member and the customer's PM coordinates with the vendor's PM to send it only the semantically compatible customer user profile. In this case, Samoa gives the vendor only those customer preferences related to shopping and reading. When the vendor receives the customer's user profile, the PM in the server host coordinates with the local SME facility to verify that the customer's user profile is semantically compatible with the vendor's discovery profile. In case of successful matching, the PM coordinates with both the PSNM to include the customer in the vendor's place-dependent social network and the GSNM to store the retrieved information in the vendor's global social-network table.

Samoa social-network extraction requires place and discovery profiles to be maintained and analyzed separately (see figure 5). This has several benefits. The manager communicates only the customer's place profile to colocated customers, thus preserving the privacy of the customer's discovery profile, which might contain confidential marketing-strategy choices. Similarly, customers distribute their user profiles only to vendors pro-

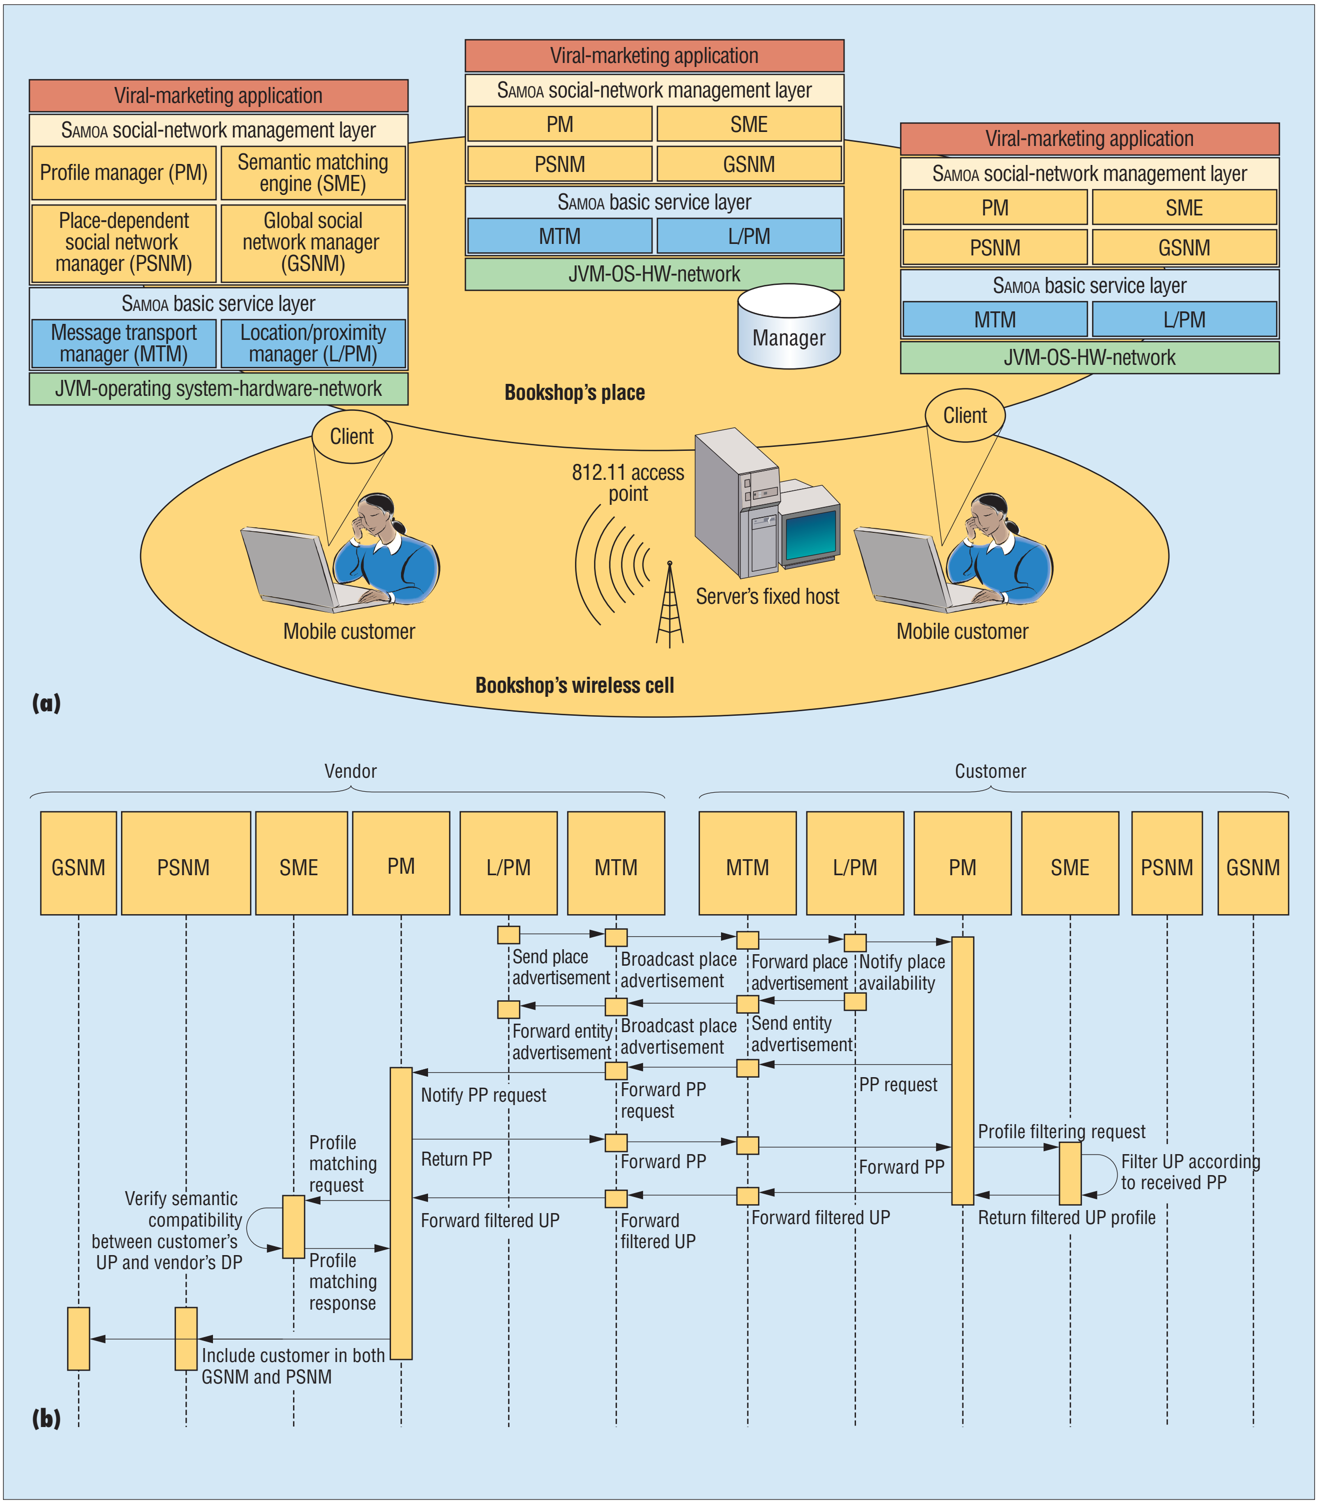

Figure 5. SAMOA facilities for social-network extraction in the viral-marketing application: (a) bookshop place abstraction and (b) interaction flow diagrams.

viding places with activities of interest. In addition, distinguishing between place and discovery profiles lets SAMOA distribute the social-network extraction overhead among all users. It performs place-profile semantic analysis only on customer devices, and semantic matching between discovery and user profiles only on vendor devices.

Finally, the vendor's place-dependent and global social networks have different uses. The application module running on the server host exploits the vendor's place-dependent social network whenever a new customer enters the bookshop. However, the visibility of the vendor's global social network lets the vendor optimize marketing strategies by browsing information about all customers who have visited the bookshop over time. For example, a vendor who detects that most customers study computer science might tailor promotions accordingly.

Performance evaluation

Using a semantic middleware, such as SAMOA, to support social-network extraction introduces several forms of overhead, depending on the deployment environment and the performance of the different middleware facilities. For brevity, we report some evaluations of the quality of SAMOA matching algorithms and the overhead that adopting semantic metadata techniques introduces. Additional implementation insights and evaluations are available at www.lia.deis.unibo.it/Research/SAMOA.

We considered a testbed activity/preference ontology modeled as a hierarchical classification tree. The tree's depth (maximum degree of activity/preference specialization) is four and its breadth (multiplicity of activity/preference related concepts) is three. Each user and discovery profile has one or two activities and two preferences for each activity, while each place profile has two characterizing activities. To evaluate our matching algorithms' quality, we measured recall—that is, the extent to which all socially compatible users are included in the network (by avoiding false negatives)—and precision—that is, the extent to which only socially compatible users are included in the network (by avoiding false positives).[14] Because our matching algorithm is complete, its recall is optimal. SAMOA has also demonstrated a good level of precision, mostly because it can look for the manager-specified preference values only in the semantically correct activity type, thus reducing false positives.

We also evaluated semantic-based matching's impact on social-network extraction time. We considered the second matching algorithm, which is the most complex. We executed the tests on an AMD Athlon XP 1600 processor, equipped with 256 Mbytes of RAM, running Windows XP Home Edition. We implemented SME using Jena 2 and Pellet 1.3 (on JDK 1.4.2). With a testbed search space of 65 eligible members, determining the manager's social network took from 8 to 9 milliseconds, depending on the manager's preference complexity. In addition, our tests showed that the most time-consuming activities are ontology parsing and querying. These activities are responsible for roughly 55 and 40 percent, respectively, of the total matching time, whereas reasoning is responsible for only about 5 percent.

Social-network creation time also depends on the communication overhead needed to transfer profiles over the wireless connection. However, the underlying wireless connection's quality and throughput heavily influence the overhead, independent of SAMOA functioning. For this reason, we don't evaluate this overhead.

The encouraging results from our early experiences with SAMOA are stimulating further research to improve the framework design. We're working along several directions, primarily on integrating SAMOA with security supports for addressing privacy issues, which is crucial to leveraging SAMOA's adoption in untrusted ubiquitous environments. ■

The Authors

Dario Bottazzi is a research fellow in the Department of Electronic and Information Systems at the University of Bologna. His research interests include middleware solutions for collaborative applications in mobile environments, social computing, and pervasive computing technologies. He received his PhD in computer science engineering from the University of Bologna. He's a member of the IEEE. Contact him at Dipartimento di Elettronica, Informatica e Sistemistica, Univ. of Bologna, Viale Risorgimento, 2-40136 Bologna, Italy; dbottazzi@deis.unibo.it

Rebecca Montanari is an associate professor of computer engineering in the Department of Electronic and Information Systems at the University of Bologna. Her research interests include policy-based networking and systems/service management, mobile agent systems, security management mechanisms, and tools in both traditional and mobile systems. She received her PhD in computer science engineering from the University of Bologna. She's a member of the IEEE and AICA. Contact her at Dipartimento di Elettronica, Informatica e Sistemistica, Univ. of Bologna, Viale Risorgimento, 2-40136 Bologna, Italy; rmontanari@deis.unibo.it

Alessandra Toninelli is a PhD student in the Department of Electronic and Information Systems at the University of Bologna. Her research interests include semantic technologies, context-aware applications, and security management in pervasive environments. She received her Laurea degree (MSc) in computer science engineering from the University of Bologna. She's a member of the IEEE and ACM. Contact her at Dipartimento di Elettronica, Informatica e Sistemistica, Univ. of Bologna, Viale Risorgimento, 2-40136 Bologna, Italy; atoninelli@deis.unibo.it.

Acknowledgments

The Italian Ministry of Education, Universities, and Research's (MIUR) Basic Research Investment Fund (FIRB) for Knowledge-Oriented Technologies for Internet-Based Enterprise Aggregation (TOCAI), the Italian National Research Council Strategic Middleware for Mobile Ad Hoc Networks and their Applications (IS-MANET), and the MIUR Project of Significant National Interest (PRIN) Middleware Approach to Mobile Multimodal Services (MOMA) projects supported this work.

References

1. S. Wasserman et al., *Social Network Analysis: Methods and Applications*, Cambridge Univ. Press, 1994.
2. M. Foth, "Facilitating Social Networking in Inner-City Neighborhoods," *Computer*, vol. 39, no. 9, 2006, pp. 44–50.
3. E.F. Churchill et al., "Social Networks and Social Networking," *IEEE Internet Computing*, vol. 9, no. 5, 2005, pp. 14–19.

4. L. Terveen et al., "Social Matching: A Framework and Research Agenda," *ACM Trans. Computer-Human Interaction*, vol. 12, no. 3, 2005, pp. 401–434.
5. G. Kortuem et al., "Wearable Communities: Augmenting Social Networks with Wearable Computers," *IEEE Pervasive Computing*, vol. 2, no. 1, 2003, pp. 71–78.
6. Q. Jones et al., "People-to-People-to-Geographical-Places: The P3 Framework for Location-Based Community Systems," *Computer Supported Cooperative Work*, vol. 13, nos. 3–4, 2004, pp. 249–282.
7. M. Harada et al., "Finding Authoritative People from the Web," *Proc. 4th ACM/IEEE-CS Joint Conf. Digital Libraries* (JCDL 04), ACM Press, 2004, pp. 306–313.
8. M. van Alstyne et al., "EmailNet: A System for Automatically Mining Social Networks from Organizational Email Communication," *Proc. 1st North Am. Assoc. Computational Social and Organizational Science Conf.*, NAACSOS, 2003, www.casos.cs.cmu.edu/events/conferences/2003/proceedings.html.
9. R. Kumar et al., "The Web and Social Networks," *Computer*, vol. 35, no. 2, 2002, pp. 32–36.
10. H. Kautz et al., "The Hidden Web," *AI Magazine*, vol. 18, no. 2, 1997, pp. 27–35.
11. P. Mika, "Flink: Semantic Web Technology for the Extraction and Analysis of Social Networks," *J. Web Semantics*, vol. 3, no. 2, 2005, pp. 211–223.
12. F. Wang et al., "Social Computing: From Social Informatics to Social Intelligence," *IEEE Intelligent Systems*, vol. 22, no. 2, pp. 79–83.
13. A. Pentland, "Socially Aware Computation and Communication," *Computer*, vol. 38, no. 3, 2005, pp. 33–40.
14. A. Bernstein et al., "Toward High-Precision Service Retrieval," *Proc. Int'l Semantic Web Conf.*, LNCS 2342, Springer, 2002, pp. 84–101.