Enhancing E-Communities with Agent-Based Systems

Intelligent software agents offer an ideal technology platform for providing data sharing, personalized services, and pooled knowledge while maintaining user privacy and promoting interaction in e-communities.

Simon Case
Nader Azarmi
BTexact Technologies

Marcus Thint
Concert Global Networks

Takeshi Ohtani
Fujitsu Laboratories

Traditionally, the term *community* refers to a location where people with common interests gather to share experiences, ask questions, or collaborate. Because they are present in the same locale, members can meet easily to learn from each other by sharing their explicit knowledge and revealing information about their successes and failures.

Alternatively, people who share a common interest, such as members of a profession, can join to form a wider community. Such communities have a deeper reservoir of knowledge, but they meet only occasionally at periodic events like conferences, which serve as forums for exchanging knowledge and ideas.

Professional communities offer obvious advantages. Members benefit from asking questions and sharing their explicit knowledge with one another. This knowledge sharing provides a more effective way of learning because it avoids duplicating the current or past efforts of other members.

Similarly, e-communities provide the advantage of connecting geographically disparate groups. These communities use Web technology as a vehicle for disseminating knowledge and information more quickly and inexpensively as well as for global communication and collaboration. Like traditional communities, e-communities act as repositories of information for their members, but they can store a larger amount of important data.

E-COMMUNITY CHALLENGES

An e-community must achieve at least two functions:

- Its members should be able to locate or be directed to relevant people and stored information.
- It should foster community spirit so that members feel they belong to a group with an identity.

An e-community site must offer services that maximize the information benefits to its members while controlling information overload, promoting member interaction, and maintaining community involvement for the benefit of all members.

E-community sites must regulate the amount of communication and information flow intelligently so that the burdens of membership do not become greater than its benefits. Members should have access to all the information they want without feeling overloaded. To foster community spirit, e-community members should be able to locate and be informed about other like-minded members, and the community should announce important community-wide news and events to interested members.

In a traditional community, selected members control the information flow either by one-on-one communication or by broadcasting to the group. Typically, for example, certain members are responsible for

0018-9162/01/$10.00 © 2001 IEEE

informing appropriate people about upcoming events. Other members, regarded as experts on certain topics, answer queries or direct individual members to useful sources of knowledge.

An e-community can include functionality that lets members disseminate information more easily to the correct people. Further, community members can use various technologies to personalize information and regulate its routing automatically.

Intelligent software agents offer an ideal technology platform for providing novel services and solutions that will meet the challenges of building successful e-communities.

INTELLIGENT SOFTWARE AGENTS

Intelligent agents have some or all of the following capabilities: cooperation, proactivity, and adaptability. Figure 1 shows a diagram of agent taxonomy. *Cooperative agents* communicate with other agents and act according to the results of that communication. *Proactive agents* initiate actions without user prompting. *Adaptive agents*, learning from past experience, change how they behave in given situations. *Personal agents* are proactive and serve individual users. *Collaborative agents* are proactive and cooperate with other agents.

Adaptive personal agents are an ideal technology for finding a user's personalized information. Researchers have developed personal software agents to help manage the increasing amount of electronic information available.[1-4] Because these agents can initiate tasks without explicit user prompting, they can undertake tasks in the background, such as searching for information. The user can access the search results by various means, but coupling them with Web technology allows easy access from wherever the user views the site.

Equally important, some agents learn from experience. In the context of e-communities, learning more about individual members facilitates updating each member's profile; over time, this improves the accuracy of community data, including information about documents, people, and contacts. Personal agents both produce and consume information. By sharing their domain's knowledge with other agents, subject to the privacy limitations imposed on them, they contribute further to community knowledge.

Collaborative filtering agents specialize in promoting interaction among community members. These agents benefit both senders and recipients because users can broadcast information to those who are interested in it without annoying other members. Contact-finding agents can locate members with distinct interests or competencies so that members can find experts in a given subdomain or other members with interests similar to their own. Agents can also work on behalf of individual members, shielding them from excess information or protecting expert members from excessive requests, thereby maintaining membership benefits for those who might otherwise become overwhelmed.

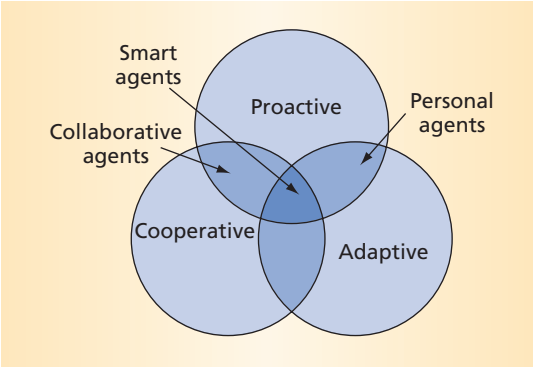

Figure 1. Agent taxonomy. The circles correspond to general agent capabilities. Smart agents exhibit a combination of all capabilities. Cooperative agents communicate with other agents; adaptive agents behave differently in given situations; personal agents serve individual users; proactive agents initiate actions; collaborative agents cooperate with other agents.

AGENT PORTFOLIO

BTexact Technologies, a division of British Telecommunications (BT), has developed a portfolio of agents consisting of a central profile-manager agent and a suite of application agents that use profiles in conjunction with several information sources. The personal agent framework (PAF) integrates several agents.

Profile manager

The e-community site hosts a *profiler* agent for each user. This agent stores user interest information in a hierarchy in which interests lower in the hierarchy inherit their parent interests' characteristics. The profile is transparent to the user, who has complete control of what it contains. The user can set each interest to be private, restricted, or fully public. In the case of interests labeled private, no other community members know that the user has this interest. On the other hand, users can openly share public or restricted interests with other community members.

Users can access the Web-based agent from any Internet connection that links to the e-community Web site, subject to an initial user login and authentication. Users can add new interests from a predefined hierarchy, which they can further edit to suit their personal preferences. Alternatively, users can create their own interest categories or they can install a client that monitors their activity at the desktop, automatically generating potential interest areas they can add to their profile.[5]

Application agents

Bugle uses profile information to generate a daily newsletter that contains articles relevant to the user's personal interests. Upon viewing the Bugle main page

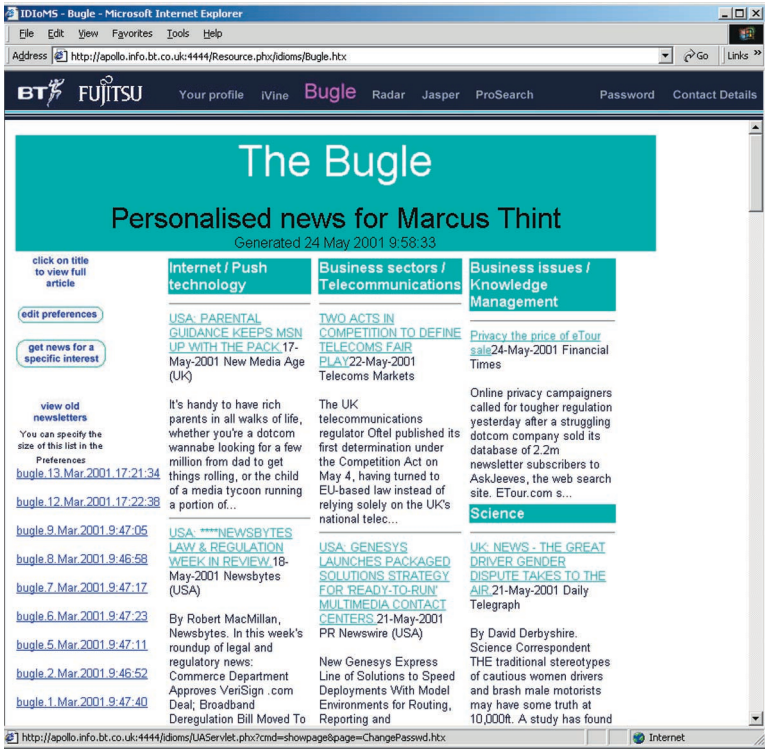

Figure 2. Sample Bugle main page. Bugle filters articles of interest and presents the user with a personalized newspaper.

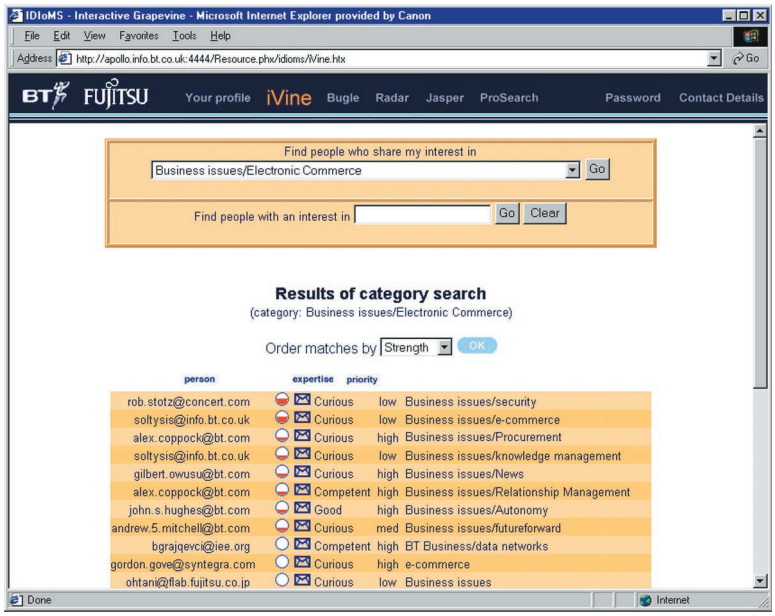

Figure 3. Sample output from iVine, which uses profiles to locate other members with common interests and displays a list sorted by strength of match or level of expertise on a particular topic.

on the e-community site, the user sees the headlines along with a summary and a link to the full article, as Figure 2 shows. A list of keywords that the agent deems important appears at the end of each full article. To semiautomatically update their profile, users can add any of these keywords under an appropriate interest category. From system trials, we found that users like to maintain control over their profiles, so the Bugle agent suggests profile updates but affects modifications only with user approval.

Grapevine is an agent that works in the background, periodically notifying members via e-mail about others who have similar interests and with whom they may wish to network.

Short for interactive-Grapevine, *iVine* lets an e-community member locate other members with similar interests. As Figure 3 shows, the member queries iVine to find other members who show interest in a specific topic. iVine provides a list of names that includes e-mail addresses and contact details, if available. iVine then shows other shared areas of interest so users can determine whether the suggested contact is appropriate for a given task. If a member has marked a given interest as private, iVine cannot access it, which helps members retain their privacy and blocks undesired communication.

Pandora, an agent that helps broaden user interests via collaborative filtering, suggests new interests for members to explore based on the interests of like-minded individuals in the community. With this kind of agent, members occasionally discover documents or activities they really enjoy and which they may not have pursued otherwise.

Users can selectively invoke *Radar,* a just-in-time information agent. As Figure 4 shows, Radar monitors a user's current activity while, for example, authoring a document, and offers relevant information resources, news reports, frequently asked questions, or useful contacts from iVine.

Community members can use *Jasper,* an agent that facilitates information sharing, to distribute a Web page or document to the community without first determining which specific members might be interested in the material. Jasper collaborates with the Profiler agent to determine who should receive the material. In addition, Jasper suggests improvements to the member's profile much like Bugle does.

ProSearch works in the background, using profiles to locate Web pages members are likely to find interesting. It removes duplicate results, pages a user has seen before, and pages that no longer exist.

An agent that other agents typically invoke, *Prosum* summarizes large documents to give recipients a précis of a particular document.

Personal assistants

In addition to having information, disseminating agents, e-communities also need to have application agents to assist with transactions between members. For members who do not want to be restricted to the electronic domain, agents that can bridge the gap between the online and physical worlds can be helpful. For example, with a *diary assistant*, members can use flexible terms such as "at the end of this week" or "in

the morning" to enter meetings. Members can also add new events while maintaining as much of the previous schedule as possible. In keeping with the philosophy of personal agents, the assistant uses personal preferences for scheduling, such as "this member prefers meetings after 10 o'clock in the morning," even when the user has not explicitly specified the preference.

Similarly, *e-mail* and *telephone assistants* learn which members a user approves being contacted by, and they use these preferences to prioritize incoming messages.

ADVANCED AGENT ARCHITECTURES

To add flexibility and scalability, we integrated PAF with open agent middleware, a Java virtual machine platform developed at Fujitsu Laboratories. We used OAM to construct flexible, dynamic, scalable, and robust distributed systems over the Internet as multiagent systems. When a network adds a new agent, OAM seamlessly and transparently arranges the agents' organization, adapts the agents within the organization to each other, and manages their collaboration. Together, PAF and OAM support sharing, managing, searching, and presenting widely distributed information across a network such as the Internet and distributed e-communities.

As Figure 5 shows, the Intelligent Distributed Information Management System (IDIoMS) automatically provides users with timely and relevant information while minimizing the need to search heterogeneous content sources for information. This enhanced system benefits service providers through the plug-and-play provision of information services, which helps providers make their services widely available.

OAM provides three main enhancements that enable dynamic flexibility and scalability.

Distributed mediation

OAM provides multiple mediator agents—federated and distributed over the network—that collaborate to service requests from applications.[6] Mediators offer a service request brokerage to other agents over the network that provides access to metainformation databases to select resources that satisfy the request. Thus, agents can access resources even though the applications are not hardwired to specific information sources. Mediators can also collate and summarize responses from the resources.

When a service provider plugs a service agent into the OAM, the agent advertises metainformation about its service to a mediator. The administrators can use XML or other formats they establish in advance for the advertisement. As Figure 6 shows, the mediator stores the advertised information in its own *condition table*, which contains condition and destination pairs, then forwards the advertisement to neighboring mediators. An HTTP-based Java servlet interfaces the resource to IDIoMS.

Mediation is based on matching a requested service with the set of services the mediator knows about.

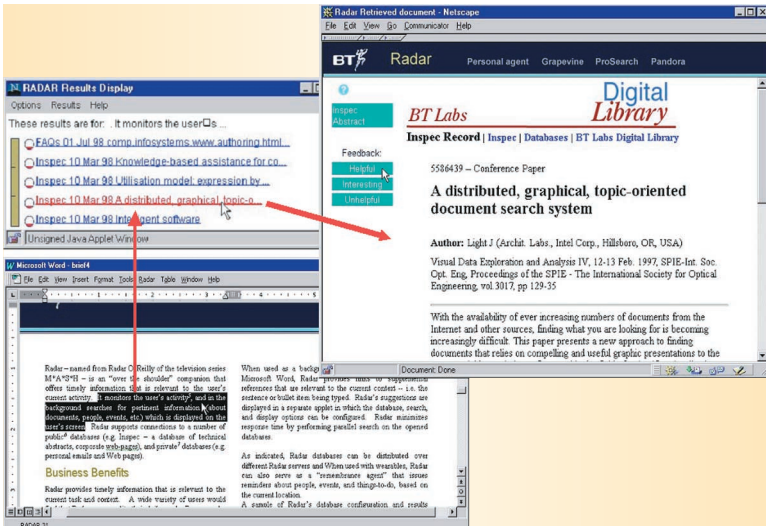

Figure 4. Radar suggests relevant references specifically for the text selected at the bottom left portion of the screen. The agent proactively searches for the information related to the current context, freeing the user to concentrate on the primary work without interruption.

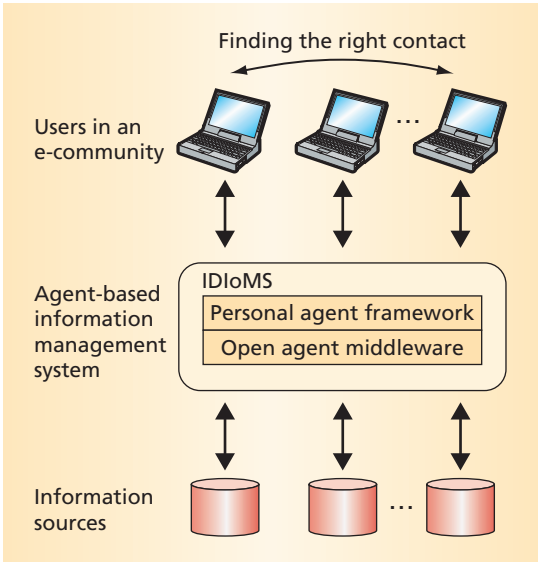

Figure 5. Intelligent Distributed Information Management System (IDIoMS) framework and application. IDIoMS provides a comprehensive set of functions for information delivery and straightforward management of distributed information resources. The scalable framework for plug-and-play content and services facilitates the use of heterogeneous content.

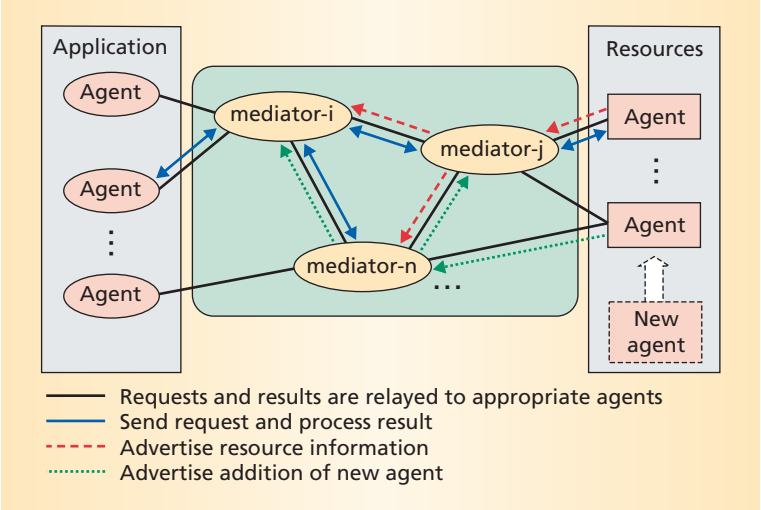

Figure 6. Distributed mediation process. The mediators accept the advertised information and send it to neighboring mediators repeatedly, as the red and green lines indicate. Thus, all mediators know the appropriate destination to route each user's request.

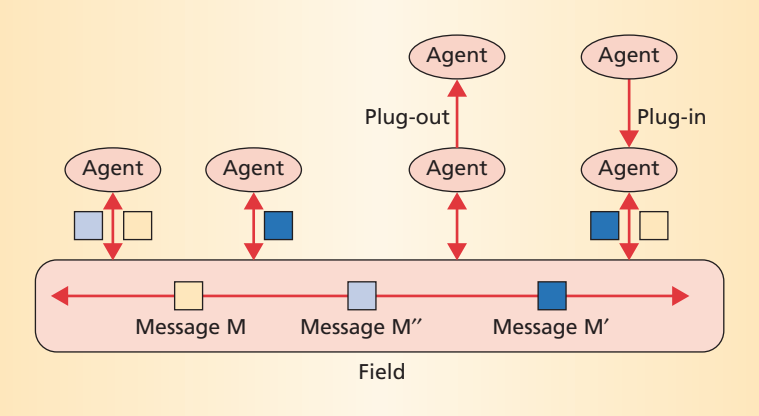

Figure 7. Communication via field. Each agent monitors the messages flowing in the field and reacts to the messages according to its own criteria. The agents are loosely coupled to the field to facilitate the addition or removal of an agent without stopping the whole system. Moreover, minimal regulation facilitates collaboration among the agents.

Upon receiving a service request, the mediator forwards it to the correct service agent or another mediator agent according to the search condition. The service agent then returns the results to the sender.

Reflective adaptation

Each service agent can potentially use different data formats and protocols when dealing with its content providers. Therefore, an organization of agents must agree on data formats and protocols before beginning collaboration. Within the OAM, agents can update interfaces and protocols for new interactions and collaborations at runtime.

Reflective adaptation is based on the agent and action programming model. In Pathwalker agent components,[7] interfaces and protocols serve as actions and form replaceable and extensible modules. Once the agent organization agrees on the interfaces and protocols, actions from repositories replace or add each agent's actions dynamically with or without collaboration with mediators. The agent that has already learned the new protocol can also provide the module.

Reflective adaptation involves changing the Java servlet code dynamically to add new features to the interface or even to update existing code—for example, adding an extra parameter to a method call.

Communication via field

Collaboration among an organization of agents requires a common medium that is flexible, dynamic, and scalable enough to enable open communication for collaboration over the Internet. Communication via field provides an event-driven, multicast-based approach to this type of collaboration.[8] When agents use a message blackboard as an information-sharing medium, the field functions as both a logical network and shared memory.

As Figure 7 shows, in communication via field, all agents on the field listen to all messages in the field, and each agent reacts to messages according to its own criteria-named patterns. Agents on the field can be added to and deleted from the field at any time independently of other agents. The agents' patterns can also change dynamically, which makes field collaboration flexible and dynamic.

USER PERSPECTIVES AND TRIALS

From the user's view, IDIoMS provides unified access to relevant and useful information. IDIoMS is built on the notion of *functional domains*—a categorization of the community IDIoMS serves. Typically, each domain is a collection of application services, including a PAF server that contains user profiles, associated agent services, and information resources. With all agent services added to the platform, IDIoMS provides a powerful resource for e-community members, both in terms of offering personalized services such as Bugle, Radar, and ProSearch and by facilitating networking and knowledge sharing with community members through services such as Grapevine, iVine, Pandora, and Jasper.

BTexact has conducted several internal trials within BT and one external trial of PAF and IDIoMS at Concert, a global venture between BT and AT&T. During these trials, one community reached more than 1,000 members. Fujitsu Laboratories has also localized some of the agents and conducted an internal trial.

Users enjoyed their daily personalized newspaper, but sometimes found that iVine matches were not fully accurate. Investigation of this problem revealed that some members did not update their contact information and interests, which highlights community-based services' dependency on members' contributions. If users choose to retain control of their profiles, they must update their records to keep the data accurate.

At first, user feedback indicated a need to control the frequency of suggestions from Grapevine and Pandora, and those agents gained acceptance after users received the option to do so. At one BT engineering center, the user community applied IDIoMS to support its collaborative work on Internet and intranet services and applications. Many other users joined the community to get assistance in performing their daily work and information-retrieval tasks. Although most users have a technical background as research scientists and software engineers, we also received favorable feedback from managers and marketing personnel. Thus, we plan to promote subsequent versions among a wider audience with diverse backgrounds.

Our trials have shown that personal agents can work effectively to serve the members of an e-community while preserving their privacy. Future research should include further investigation of several topics, including automatic learning of short- and long-term user interests, improving information retrieval accuracy, and ontology management. Clearly, successful deployment of multi-IDIoMS systems requires addressing ontology management—an issue that generally plagues the Internet. However, recognizing and efficiently collating information about similar topics within heterogeneous data sources remains a difficult issue to resolve.

Using XML and other standards can mitigate some concerns, but it does not provide a solution to other key issues, especially with regard to legacy data. Nonetheless, we believe that IDIoMS offers an advanced agent-based information management system that provides a powerful application for use in complex, distributed Internet e-communities. ✸

References

1. P. Maes, "Agents that Reduce Work and Information Overload," *Comm. ACM*, vol. 37, no. 7, 1994, pp. 31-40.
2. Y. Lashkari, M. Metral, and P. Maes, "Collaborative Interface Agents," *Proc. 12th Nat'l Conf. Artificial Intelligence*, AAAI Press, Menlo Park, Calif., 1994, pp. 444-450.
3. H. Nwana, "Software Agents: An Overview," *The Knowledge Eng. Rev.*, vol. 11, no. 3, 1996, pp. 205-244.
4. B. Crabtree, S. Soltysiak, and M. Thint, "Adaptive Personal Agents," *Personal Technologies J.*, vol. 2, no. 3, 1998, pp. 141-151.
5. S. Soltysiak and B. Crabtree, "Automatic Learning of User Profiles—Towards the Personalisation of Agent Services," *BT Technology J.*, July 1998, pp. 110-117.
6. T. Mohri and Y. Takada, "Virtual Integration of Distributed Database by Multiple Agents," *Proc. 1st Int'l Conf. Discovery Science* (DS 98), Lecture Notes in Computer Science, vol. 1532, Springer-Verlag, Heidelberg, Germany, 1998, pp. 413-414.
7. I. Iida et al., "Multi-Agent Platform for Seamless Personal Communications," *Telecom 99 Forum Interactive Summit*, Int. 7, ITU Publications, Geneva, 1999; http://www.itu.int/TELECOM/wt99/cfp/auth/2644/pap_2644.doc.
8. T. Iwao et al., "Flexible Multi-Agent Collaboration Using Pattern Directed Message Collaboration of Field Reactor Model," *Proc. 2nd Pacific Rim Int'l Workshop on Multi-Agents* (PRIMA 99), Lecture Notes in Artificial Intelligence, vol. 1733, 1999, Springer-Verlag, Heidelberg, Germany, pp. 1-15.

Simon Case *is a senior research scientist at Intelligent Systems Lab, BTexact Technologies, Ipswich, UK. His research interests include fuzzy and probabilistic learning and personal agent systems. Case received a PhD in artificial intelligence from the University of Bristol. Contact him at simon.case@bt.com.*

Nader Azarmi *is a senior technology manager at Intelligent Systems Lab, BTexact Technologies, Ipswich, UK. His research interests include agent technology, planning and scheduling, and soft computing. Azarmi received a DMPhil in computer science from Essex University. Contact him at nader.azarmi@bt.com.*

Marcus Thint *is a principal member of the technical staff at Research and Applications Group, Concert Global Networks, Reston, Va. His research interests include a soft computing approach to machine learning, pattern recognition, and robotics. Thint received a PhD in machine intelligence and robotics from Duke University. Contact him at marcus.thint@concert.com.*

Takeshi Ohtani *is a researcher at IP Server Project Group, Fujitsu Laboratories Ltd., Fukuoka, Japan. His research interests include resource discovery, personal agents, and multiagent systems. Ohtani received an MSc in information engineering from Tohoku University. Contact him at ohtani@flab.fujitsu.co.jp.*