The Requirements for Personal Mobile Assistants in a Mobile Telecommunications Environment

Barcin Kozbe and Michael N. Huhns
Center for Information Technology
Department of Electrical and Computer Engineering
University of South Carolina
Columbia, SC, U.S.A. 29208
(803) 777-5921 or {kozbe,huhns}@ece.sc.edu
(803) 777-8045 FAX

Abstract

The telecommunications environment that will exist in the near future will consist of a complex web of interacting devices and services. Users will need help in dealing successfully and productively with this environment, which will be made more complex by their own mobility. They will need on-line software systems that can represent them in their absence, represent them in routine or mundane interactions, and assist them in complex interactions. We term such on-line, software systems "assistants." The complexity, dynamic, nature, and openness of the environment will necessitate intelligence for the assistants. The unique demands and needs of each user will require assistants that are customized or personalized. The proliferation of agents will require that they be able to interact with each other productively. In this paper we consolidate these requirements into an architecture for interacting personal assistants.

1 Introduction

Mobility/mobile and agent are terms that are commonly used by computer scientists and communications professionals. The ability to communicate whenever and wherever one is, and to access large amounts of distributed information sources located anywhere and at anytime, are needed to meet the requirements of today's changing social and business world. A common definition of a personal communication service is "the ability for a user to originate and receive calls or messages anytime, anywhere, and in several roles using a pocket terminal" [1]. People equipped with portable-mobile devices, such as laptop computers, PDAs, mobile phones, and pagers, should be able to move and work freely among different locations. Having a personal number associated with communication and information services will be one of the technological changes that will integrate our activities and make us more effective. However, in addition to this service, people need personalized assistants that retrieve and filter information, schedule meetings, book tickets, buy and sell items, etc. on their behalf. In this paper, we define Personal Mobile Assistants that perform such tasks for a user, as indicated in Figure 1.

2 Background and Analysis

The requirements of large, complex, dynamic, distributed communication and computation environments have motivated much of the research on intelligent agents. In recent years, applications of increasingly intelligent agents have appeared in telecommunications, electronic commerce, industry, simulation, information retrieval and filtering, and so on. Many agent architectures have been proposed for these applications, including interface agents, softbots, mobile agents, information agents, and believable agents. In general, an *intelligent agent* can be defined as a software program that performs some tasks in an environment with autonomous processing capabilities, so “any property of an agent must therefore be defined in terms of the task and the environment in which the task is to be performed” [2]. When thought of as “assistants,” agents may be classified as local agents and mobile agents.

Local Agents These agents, like interface agents [3], perform different kinds of tasks as they collaborate with their local user or other agents: information filtering and retrieving, scheduling meetings, mail management, etc. These agents act as Personal Assistants, and they can learn users' habits and preferences, and teach other users or agents.

Mobile Agents A new paradigm in agent technology is mobility. Mobile agents move through a network from one place to another with state information and executable code. This paradigm is called remote programming (RP), and is an alternative to remote procedure calling (RPC) [4]. During their travel, mobile agents may meet other mobile or static agents, e.g., at virtual markets. Application areas for mobile agents are network management simulation, information retrieval and filtering, mobile computing, and

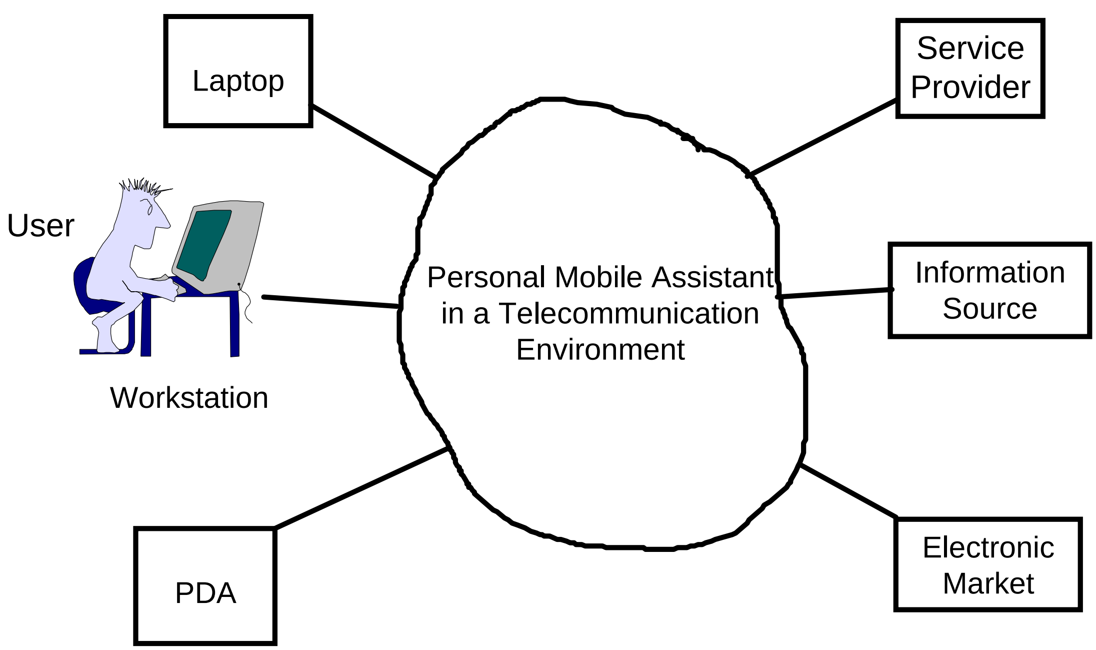

Figure 1. The telecommunication environment for a user, mediated by a personal mobile assistant

electronic marketplaces.

2.1 Personal Mobile Assistants

Via the Internet, one can currently gain access to a large number of distributed information sources and electronic markets to retrieve information that is wanted, or buy and sell items one needs. But in an open, dynamically changing environment such as the Internet, users need assistance, and this should be in the form of mobile intelligent agents. We term such agents Personal Mobile Assistants (PMA). To aid in understanding the role of a PMA, we give an example. Offices typically have different kinds of assistants for different kinds of tasks. There is a secretarial assistant who shares the same work environment as the rest of the office workers and answers incoming calls, schedules meetings, informs them regarding important events, etc. But many offices also have a personal assistant who goes to office-supply stores or bookstores to buy needed items. A PMA is similar to the latter kind of assistant.

A PMA is distinguished from personal agents by its mobility, and from mobile agents by its task and environment. A mobile agent is a very general concept within the framework of a task and its associated environment in which the agent executes applications. Before we define the properties of a PMA, we will describe its capabilities and environment, since the properties of an agent depend on its environment and task.

A task is a description of what the agent is supposed to achieve in the environment [2]. Possible tasks for a PMA include information retrieval and filtering, buying and selling items, making passive documents appear active, keeping track of its user, and managing interactions with other agents.

The environment in which a PMA operates is dynamically changing and nondeterministic. Because the environment is complex, it is difficult to keep track of all of its inaccessible aspects. Based on the task and the environment, the three main properties of a PMA are agency, intelligence, and mobility.

2.2 Agency

Agency defines the basic characteristics of PMAs. It has the following descriptive facets:

Autonomy. An important characteristic of PMA agency is *autonomy*, i.e., a PMA is an autonomous software entity. It can take initiative and operate in a dynamic environment without interacting with a user. This is one of the major features that distinguishes a PMA from other software programs. A PMA also has some degree of control over its own actions.

Social Ability. A PMA can interact with other agents. The interaction can appear in the form of cooperation or competition. A PMA might cooperate with others; for example, at a meeting point or electronic market, it can find out via communication that its interest in some kind of information is similar to that of other PMAs. They can then collectively gather the information, thus sharing in the cost. A PMA might be competitive when selling or buying some goods. To compete effectively, it might make use of a utility function that helps it determine the best strategy to pursue.

Adaptability. A PMA has an ability to change its behavior to be able to reach its goal despite environmental changes.

Reactivity. A PMA responds to changes in its environment, and an agent is reactive if it is able to respond in a timely manner to such changes.

Communication Ability. Explicit communication is needed among PMAs to enable them to cooperate. A PMA typically communicates with others via a high-level communication protocol, such as KQML.

2.3 Intelligence

Intelligence introduces a stronger notion of agency. It gives a PMA the ability to learn or understand from experience and make decisions in new situations without interacting with a user.

2.4 Mobility

A PMA moves through a network from one place to another, entering remote computers and executing applications. This paradigm brings benefits as a cost and time saving.

Of course, it is possible to add other properties that extend the character of a PMA. But we consider the attributes we listed here to reflect common and basic features of a PMA.

2.5 Electronic Commerce

Increased demands for time and cost savings and the advent of new technologies in communications and computers are leading to changes in the way products are introduced and business is conducted—this is being termed “electronic commerce.” The term also includes information search and retrieval when used to assist people in making decisions on buying or selling items. There are three communication patterns in electronic commerce: (1) person-to-person, (2) person-to-application, and (3) application-to-application. Telephone, fax, and email are used in person-to-person communications, the World Wide Web is used for person-to-application communications, and electronic data interchange (EDI) is used for application-to-application communications.

On-line services, on-line marketing, electronic purchase orders, shipping information, and electronic payment are appearing in standardized forms under EDI, and EDI servers will soon be complementing Web servers [5]. This will automate many of the routine aspects of how business is conducted in the world. Currently there are several on-line shopping services, such as Bargain Finder of Andersen Consulting, and electronic marketplaces, such as America On-line and CompuServe. Producers are offering a wide array of products and services to global consumers via the Internet, and consumers can find the cheapest and the best offerings among these services. However, information overload and the huge number of available services make it difficult for a user to find desired information or products effectively. In this case, there is a need for an intelligent mobile agent that assists its owner while he takes care of other activities. In electronic commerce, a PMA can act as a buyer by finding the best offerings that meet its owner’s requirements, or as a seller by providing service and products or advice for its user.

3 PMA System Architecture

There are many problems in developing a full system architecture for a PMA that meets the requirements defined in Section 2. System complexity is increased with personal mobility, and there is a need for a set of communication protocols and information access mechanisms to determine the location of a user. There has been a great deal of research and development done in both the communication and computing areas that supports the creation of an environment for Intelligent Mobile Agents. Some of the major topics are the following:

- Oracle's Mobile Agents supporting Cellular Digital Packet Data (CDPD)
- Mobile Agents and Telescript [4]
- Itinerant Agents for Mobile Computing [6]
- Smart Networks and Intelligent Agents [7]
- Negotiation [8]
- Mobile Computing
- The Universal Mobile Telecommunication System (UMTS)
- Global Systems for Mobile Communication (GSM)
- Research on other topics in Distributed Artificial Intelligence (DAI) and Multiagent Systems (MAS)

However, we are aware that problems remain, and open questions for future investigation involve:

- Communication languages and network protocols supporting distributed mobile computing.
- Security aspects in mobile communication and computing
- Management of mobility
- Resolving semantic differences among heterogeneous knowledge resources
- Handling information lifetime
- Goal-based behavior

Our research will focus on the goal-based behavior of a PMA.

4 Goal-Based Behavior

A default goal for a PMA is to improve its ability to find the best partner, buyer, or seller, or to retrieve the most relevant information. In the environment in which the PMA operates, other agents with different goals and characteristics are also operating. Some of these agents may be less reliable than others, and an agent can even lie, cheat, or misinform another agent. Similar relations can also be seen in human society, but they are dealt with by relying on trusted associates and personally verified information. We have best friends, favorite restaurants, or favorite stores. A PMA can categorize other agents according to most reliable, reliable, less reliable, or unreliable as a result of its own experience or learning from its user. However, a user is mostly unfamiliar with the environment in which the agent acts because of its large size and complexity, and he has limited control on the agent because of the agent’s mobility. The agent must therefore be controlled by its task and its environment, and this determines what the agent should model or learn about others.

A PMA needs models of itself, other agents, its user, and the environment to construct possible future scenarios that it can then use to guide its actions, as shown in our architecture for an agent in Figure 2. We assume that a PMA begins to act in an unknown environment without any initial knowledge about the characteristics of its user and other agent(s). It acquires this dynamically, and in this way is self-initializing.

4.1 User Modeling

To model its user, the PMA keeps track of him and learn his habits, preferences, and interests over time. Just as a child observes the reactions of its parents to its own actions and builds a model of correct behavior, a PMA can know how successful it is by observing its user's feedback and reactions. This parameter of the user model is employed to measure the performance of a PMA's action. A user's feedback can be positive as "well done job," negative as "too bad," or "no comment." A problem here is how a PMA might interpret this "no comment." One way is that it does not take this into consideration. But in the long term, repeated "no comments" create uncertainty in a PMA's making decision process. We consider a better approach to be that the PMA itself decides how to handle this reaction using the information in its performance system or by searching and finding similar cases from its own past experience.

4.2 Environment Modeling

Acting upon its environment and receiving information via percepts provided by the environment, a PMA creates a model of the environment in the form of an internal map. A utility function is then used to rank the states of the environment to which the PMA's actions will lead. Feedback from the environment makes the PMA more independent and can play an important role when the PMA is making decisions autonomously (i.e., without a user's feedback).

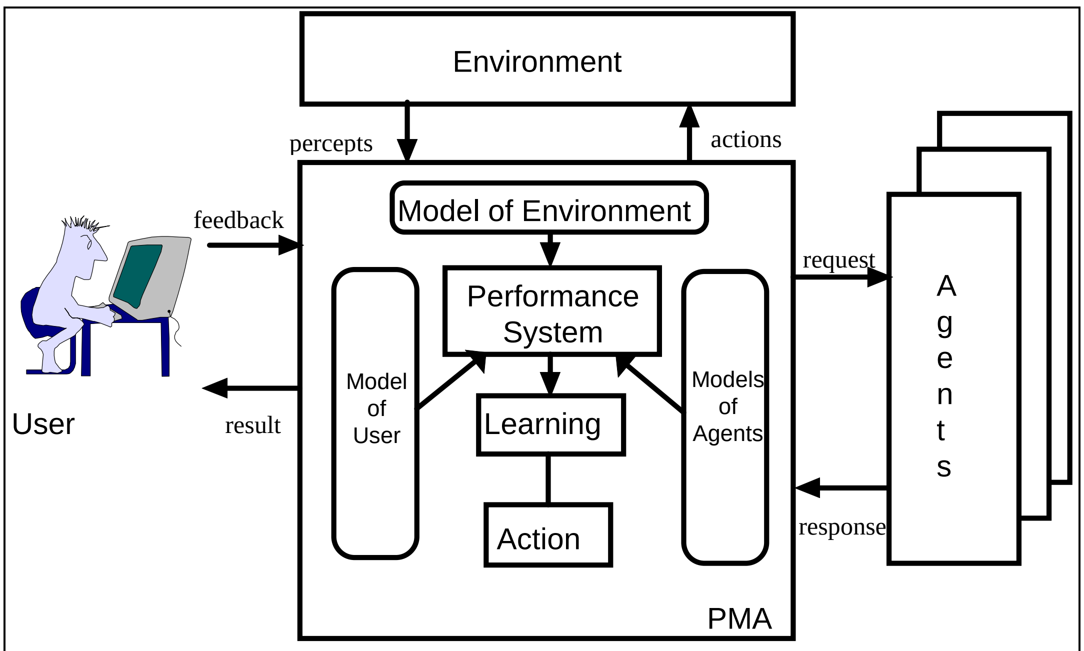

Figure 2. A PMA must learn in order to be effective.

4.3 Agent Modeling

How should one agent represent another, and how should it acquire the information it needs to construct a model in that representation? This has, we believe, a simple and elegant answer: *the agent should presume that unknown agents are like itself, and it should choose to represent them as it does itself.* As it learns more about them, it only has to encode any differences that it discovers. This can make the resultant representation concise and efficient. Some of the other advantages of this are the following:

- An agent has a head start in constructing a model for an unknown or just-encountered agent.
- An agent has to manage only one kind of model and one kind of representation.
- The same inference mechanisms that are used to reason about its own behavior can reason about the behaviors of other agents; an agent trying to predict what another will do has only to imagine what it itself would do in a similar situation.
- As information about other agents is acquired through subsequent observations and interactions, models of the agents can be updated, and will then diverge from the default.

One of the ramifications of this is that an agent constructed with a belief-desire-intention (BDI) architecture would attribute intentions to another agent, even if that other agent does not have a BDI architecture or even any explicit intentions! This is appropriate, because it makes the other agents easier to think about, analyze, and predict. This is also likely why people tend to anthropomorphize inanimate objects, i.e., to imbue them with human-like characteristics—it makes the world more uniform (everything is animate) and therefore easier to deal with.

How an agent can model other agent(s) depends on the agent architecture, the relationship between the agents, and the type of knowledge that will be used in modeling. Modeling may take place by observing another agent's actions—we call this *passive observation*—or by communicating with other agent(s) and acquiring knowledge about them, which we call *active observation*.

A PMA via an ASK-TELL or REQUEST-RESPONSE protocol may acquire and share other agents' knowledge, including their capabilities and goals. It classifies their responses, advice, and reactions by comparing the feedback from the user and environment for the action taken. This classification forms the model of other agents. Agents can be given a trust value classification of reliable or unreliable, as was done in RAD [9]. But this categorization was found to be too coarse. Instead of this, we will apply a "trust function" that makes a more flexible and fuzzy classification. We will also investigate how a PMA might detect that a less reliable or unreliable agent has improved its behavior. This feature is important to give a "second chance" to unreliable agents and produce a more realistic classification.

In information retrieval, the quality of retrieved information is commonly rated by two parameters: precision and recall. *Precision* is the ratio of relevant information retrieved to the total amount of information retrieved. *Recall* is the ratio of relevant information retrieved to the total amount of relevant information available. An information retrieval system should try to maximize both of these measures, but that causes the following conflict: one way to maximize recall is to retrieve all information, but this would result in a low value for precision, because most of the retrieved information would be irrelevant; one way to maximize precision is to return no information, but this would result in a low value for recall.

Two other important parameters involve what and where, i.e., *where* desired information is located and *what* is the content at that location. A good agent in cooperative, agent-based information retrieval will provide correct and accurate advice about both of these parameters. A simple measure of such advice should produce the following ordering (from best to worst):

1. Precise and correct information
2. Approximate information
3. Pointers to precise and correct information
4. Pointers to approximate information

Simplistically, we can rate the value of an agent's advice, and therefore the agent itself, by the product of a measure of the quality of where with a measure of the quality of what. Each of the two measures of quality would themselves be measured in terms of precision and recall.

Our research goal is to create a PMA that observes its user, other agent(s), itself, and its environment, and learns models of their dynamically changing behavior to offer an efficient, reliable, and secure service for its user.

5 Conclusion

Mobility is a key requirement for personalized services and assistants. In this paper we introduced the concept of a Personal Mobile Assistant, which differs from other definitions of intelligent agents in its capabilities. We have described the basic requirements of a PMA and examined electronic commerce as an application area. In the last section, we presented some of the major research underway and the remaining problems. Our future work will investigate the goal-based behavior of PMAs and the models of agents and environments that are necessary for successful behavior.

References

[1] Philippe Dupuis, "A European View on the Transition Path towards Advanced Mobile Communications Systems," *IEEE Personal Communication Magazine*, 1995.

[2] R. Goodwin, "Formalizing Properties of Agents," Technical Report CMU-CS-93-159, Department of Computer Science, Carnegie-Mellon University, 1993.

[3] Pattie Maes, "Agents that Reduce Work and Information Overload," *CACM*, vol. 7, 1994, pp. 31-40.

[4] James White, "Mobile Agents," White Paper, General Magic, 1995.

[5] Peter Benson, "EDI and the Internet," *Internet Systems*, May 1996, pp. 29-32.

[6] David Chess, Benjamin Grosof, Colin Harrison, David Levine, Colin Parris, and Gene Tsudik, "Itinerant Agents for Mobile Computing," IBM T.J. Watson Research Center, New York, 1995

[7] Colin G. Harrison, "Smart Networks and Intelligent Agents," IBM T.J. Watson Research Center, New York, 1995.

[8] Jeffrey S. Rosenschein and Gilad Zlotkin, *Rules of Encounter*, MIT Press, 1994.

[9] N. Arni, D. Bridgeland, J. Christian, M. Huhns, C Petrie, E. Rich, J. Shea, and M. Singh “Overview of RAD: A Hybrid and Distributed Reasoning Tool,” MCC Technical Report Number ACT-RA-098-90, Microelectronics and Computer Technology Corporation, Austin, TX, February 1990.