

NRC Publications Archive Archives des publications du CNRC

On Establishing Multi-Sensory Multi-Channel Communications Among Networked Users

Ferguson, Innes; Davlourous, J.D.

This publication could be one of several versions: author's original, accepted manuscript or the publisher's version. /
La version de cette publication peut être l'une des suivantes : la version prépublication de l'auteur, la version
acceptée du manuscrit ou la version de l'éditeur.

Publisher's version / Version de l'éditeur:

*Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI),
1995*

NRC Publications Archive Record / Notice des Archives des publications du CNRC :

<https://nrc-publications.canada.ca/eng/view/object/?id=46e22eb9-3e51-4c82-9bce-a0a6194b00cd>

<https://publications-cnrc.canada.ca/fra/voir/objet/?id=46e22eb9-3e51-4c82-9bce-a0a6194b00cd>

Access and use of this website and the material on it are subject to the Terms and Conditions set forth at

<https://nrc-publications.canada.ca/eng/copyright>

READ THESE TERMS AND CONDITIONS CAREFULLY BEFORE USING THIS WEBSITE.

L'accès à ce site Web et l'utilisation de son contenu sont assujettis aux conditions présentées dans le site

<https://publications-cnrc.canada.ca/fra/droits>

LISEZ CES CONDITIONS ATTENTIVEMENT AVANT D'UTILISER CE SITE WEB.

Questions? Contact the NRC Publications Archive team at

PublicationsArchive-ArchivesPublications@nrc-cnrc.gc.ca. If you wish to email the authors directly, please see the
first page of the publication for their contact information.

Vous avez des questions? Nous pouvons vous aider. Pour communiquer directement avec un auteur, consultez la
première page de la revue dans laquelle son article a été publié afin de trouver ses coordonnées. Si vous n'arrivez
pas à les repérer, communiquez avec nous à PublicationsArchive-ArchivesPublications@nrc-cnrc.gc.ca.

On Establishing Multi-sensory Multi-channel Communications Among Networked Users

Innes A. Ferguson and James D. Davlourous

Knowledge Systems Laboratory (IIT)
National Research Council
Ottawa, ON, Canada

innes@ai.iit.nrc.ca; james@cais.dsm.sp-agency.ca

Abstract

The PeopleFinder is a knowledge-based tool to assist users in determining the whereabouts of other staff located in an office or network environment. The tool makes use of several modes of input and output, as well as employing a number of interface and communications media with which to present information and interconnect remote system users. An accompanying video contains example uses of the tool which help illustrate some of its functionality.

Introduction

The PeopleFinder is a knowledge-based tool to assist users in determining the whereabouts of other staff located in a departmental or office environment. The tool makes use of several modes of input and output, as well as employing a number of interface and communications media with which to present information and interconnect geographically distributed system users. The People Finder employs agent-based design techniques [Fer92] as a way of integrating a variety of heterogeneous applications running on a number of different computing platforms and operating systems. Specifically, the PeopleFinder's functionality is carried out by a collection of coordinated software agents which can:

- (i) assume user-delegated tasks (e.g. interpret spoken commands, invoke email and voice recording tools, dial out on a telephone);
- (ii) perform a number of application level tasks (e.g. update the graphical user interface, apply knowledge-based heuristics to decide whether and/or how system users want to be contacted); and
- (iii) carry out a number of lower level operating system level activities, the majority of which will be beyond the level of expertise of most casual system users (e.g. translate between different speech encoding formats, establish which users are logged on to the various multi-platform computer networks throughout the building).

The remainder of this paper will provide a brief description of the main design features and functional capabilities

of the PeopleFinder. In addition to relating the PeopleFinder to some existing multimodal and/or multiagent communications applications, some comments on preliminary evaluations of the tool as well as plans for future work on the tool will be given.

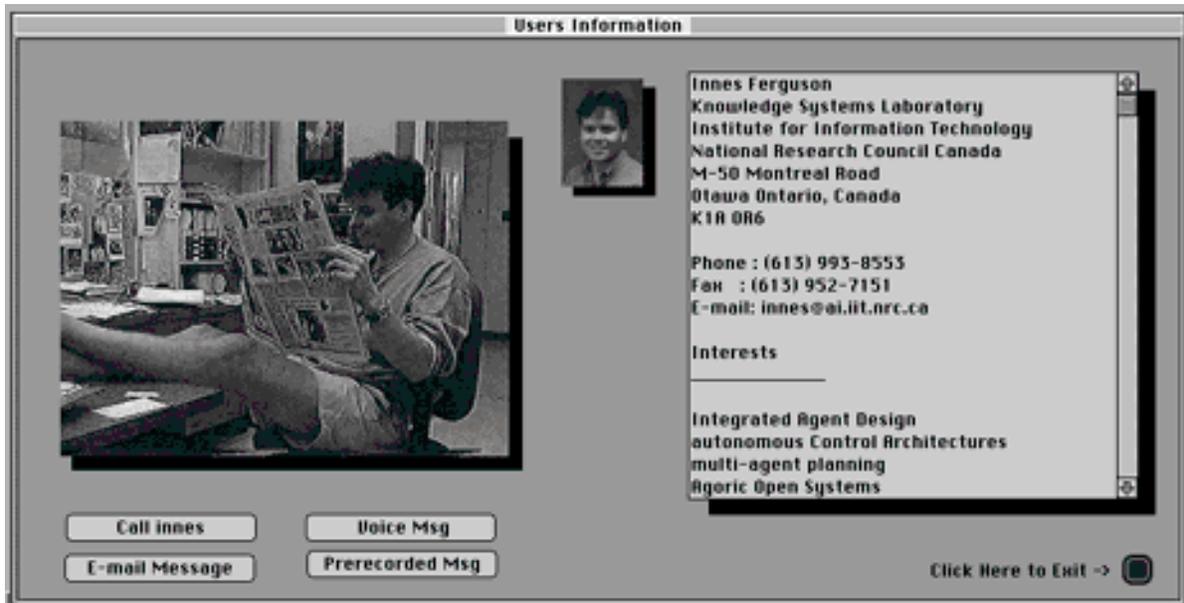
The PeopleFinder

The PeopleFinder is a system that aims to provide support for general awareness in a distributed work group. By collecting data from a number of on-line sources (e.g. login and active process statistics, electronic calendar data, telephony status information) and applying a series of heuristic rules that are designed to encapsulate typical work patterns of fellow users, the PeopleFinder attempts to inform users, as quickly and as accurately as possible, of the activities and whereabouts of other system users.

Some of the key design features of the PeopleFinder include:

- **Agent-based design.** The system is based on a number of coordinated autonomous agents, each specialized in one or more system tasks, a number of which have been mentioned above.
- **Multimodal interaction.** The system makes use of various input and output modalities for human-computer interaction: in particular, the keyboard, mouse, and speech for input; and audio and (visual) screen-based feedback for output. Figures 1 and 2, for example, illustrate some of the feedback users receive when attempting to contact other network users.
- **Multimedia presentation.** The system also makes use of a number of different media for human-computer interaction; in particular text, graphics, animation, and prerecorded video (the latter acting as a "cheap", but nevertheless very effective, substitute for in-office video cameras and/or ubiquitous computing devices such as active badges¹).
- **Communications-oriented capabilities.** The system

¹Such as those used in various in-house applications at Xerox's Palo Alto Research Center [WHFG92, Wei93].

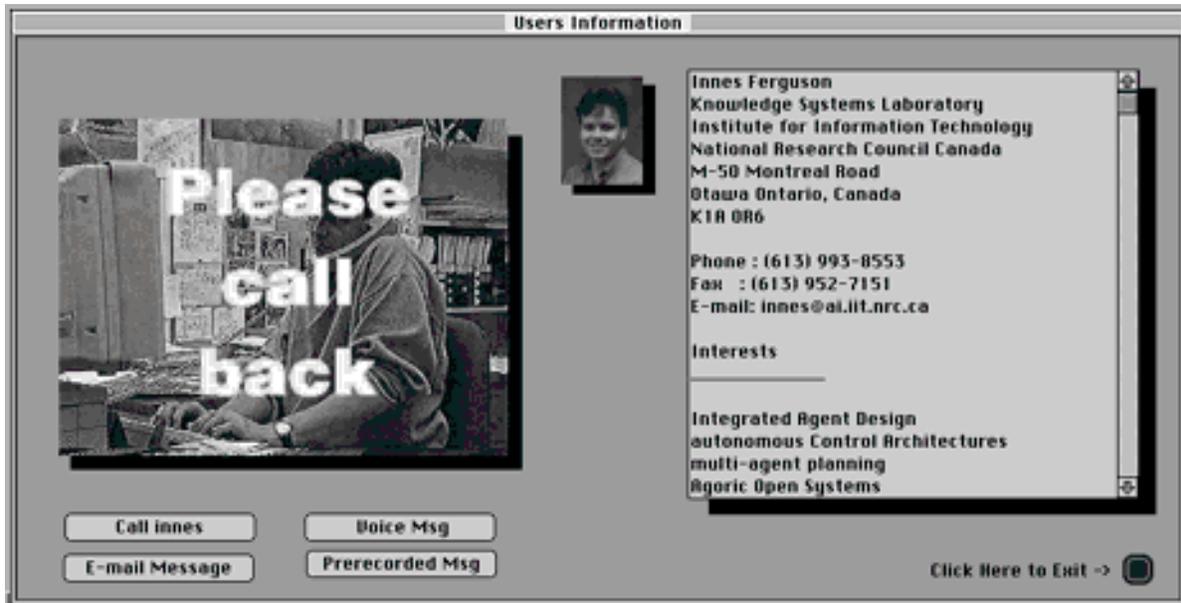


(a)

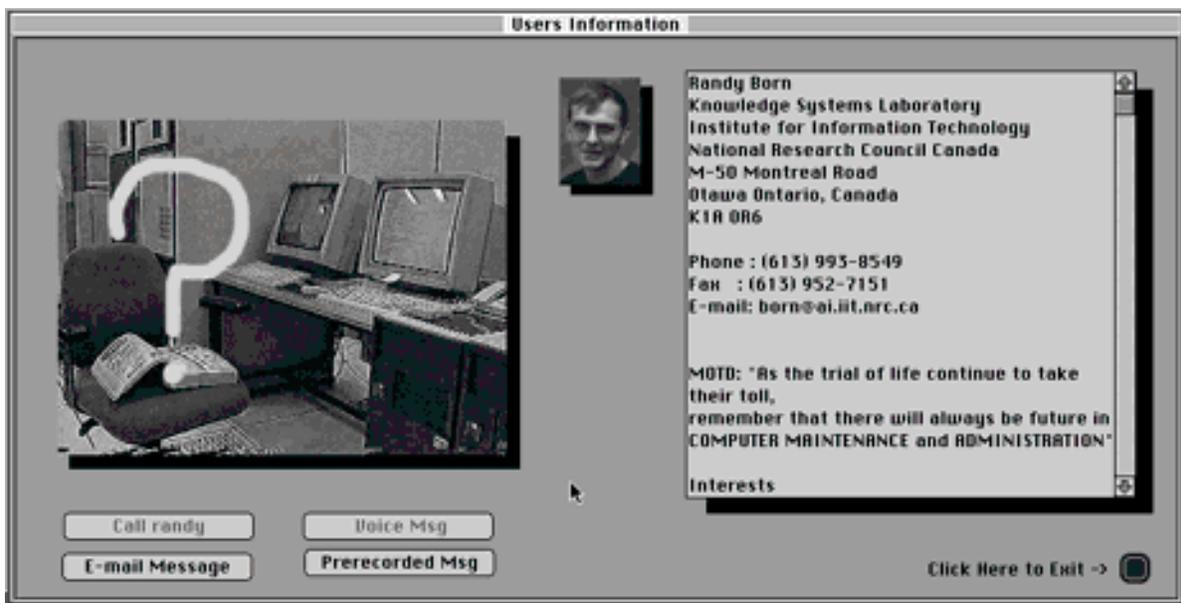


(b)

Figure 1. Example Users Information windows showing one user to be “idle” or willing to be contacted (a) and one user who is presumed to be out of the building (b). Appropriate means of communicating with the sought user are suggested in the bottom left-hand corner of each window. General contact information (e.g. information typically found in users’ “.plan” files) is listed in the scrollable text field on the right-hand side of the window.



(a)



(b)

Figure 2. More Users Information windows showing one user to be busy on the phone (a) and yet another user to be out of his office, but presumably still somewhere in the building (b). In the former case, the remote user's status is determined by placing a telephone call and monitoring the local modem for a received busy signal; in the latter case, the remote user's status is determined heuristically on the basis of elapsed idle computing time.

enables transparent communication across different computer platforms (Macintosh, Unix) and facilitates the interconnection of system users via telephone, email, and voice messaging (see Figure 3).

The agents used in the PeopleFinder application are based on the CALVIN open agent framework (see Figure 4) which in turn is an extension of the TouringMachine architecture [Fer92]. This framework provides application developers with a powerful set of agent programming tools including libraries of intra- and inter-agent protocols (e.g. KQML²), sensory and effectory apparatus, internal behavior APIs, persistent storage management, and (currently under consideration) CORBA³ compliance.

The tool is implemented using a variety of different scripting languages (AppleScript, Quickeys, and C-shell) and runs on a Macintosh Quadra 840 AV. The tool also makes use of Quicktime movies for displaying prerecorded video as well as the Macintosh's Apple Phone tool and Geoport Telecom Adapter for performing its various computer-telephony integration tasks.

Evaluation

A number of features of the CALVIN⁴ architecture have proven useful for developing multimodal applications that integrate a number of distributed media resources. In particular, rapid responses to users' commands are facilitated through integration of appropriate reactive behaviors in the system's Interface and User agents (see Figure 3); in addition, blending of complementary input modalities is facilitated through the execution of multiple concurrent agents (which in turn are able to execute multiple concurrent, task-specific behaviors).

Early evaluations of the PeopleFinder application are generally quite positive. Feedback from users of the tool within our own laboratory (which comprises between 25 to 30 staff) has enabled us to identify a list of the most popular capabilities of the tool. These include the ability to integrate a number of commonly used tools and devices (e.g. electronic mail, telephone, and "finger"-like OS-level commands); the aesthetic use of multiple media (in particular, the prerecorded video); and the ability to offer multi-sensory and multi-channel means of interaction between human and computer. In terms of negative user feedback, this was mostly directed at the less than satisfactory quality of the Macintosh Quadra 840 AV speech processing (automatic recognition and text-to-

²The Knowledge Query and Manipulation Language (KQML) is a protocol intended to support interoperability among intelligent agents in distributed applications [FMF92, Wer94].

³The Object Management Group's Common Object Request Broker Architecture (CORBA).

⁴The Communicating Agents Living Vicariously In Networks (CALVIN) architecture is an agent framework developed at the National Research Council's Knowledge Systems Laboratory (see WWW page <http://ai.iit.nrc.ca/software.html> for more details).

speech generation) capabilities.

Related Work

Besides permitting the relatively seamless integration of both non-deliberative (e.g. reactive) and deliberative (e.g. planned) behaviors, as well as providing an added level of flexibility and robustness to the overall system performance [Fer95], the existence of multiple autonomous — and, therefore, concurrently operating — agents also facilitates the application of multiple channels of interaction between user and system [FJ94]. A number of similar benefits and issues have been identified in the application of the Open Agent Architecture to the management of email; in particular, the roles of such agent skills as *delegation* ("the ability to receive a task to be performed without the user's having to state all the details"), *data-directed execution* ("the ability to monitor local or remote events, such as database updates, OS, or network activities"), and *communication* ("the ability to enlist other agents... in order to accomplish a task") [CCWB94, page 1]. In addition, the PeopleFinder can be seen to offer a number of desirable features which are characteristic of technology for *telepresence*; that is, technology which provides, despite geographical or temporal distance, a sense of social proximity — a kind of *social prosthesis* for overcoming gaps and weaknesses with an organizational structure [Luc93, Bux94].

Summary

The PeopleFinder is an application of technology which aims to support a sense of social proximity among geographically distributed users. Besides supporting presentation of diverse sources of multimedia information, the PeopleFinder emphasizes (and exploits) the use of different human sensory modalities and channels of communication in order to maximize its potential usefulness. A primary goal in designing and deploying the tool is to provide a computational system which closely matches the skills and work practices of its intended users. Early evaluations of the tool suggest the present work is heading in the right direction.

Future work already underway includes porting the graphical user interface portion of the PeopleFinder to run on both PC and Unix platforms (in the interest of extending the tool's audience and ensuring a more thorough testing and empirical evaluation phase of the project⁵); integrating a number of other software applications such as teleconferencing, voice dictation, and video camera-based face recognition; extending agents' capabilities for autonomously resolving run-time conflicts

⁵In fact, we plan to make the PeopleFinder available on the WWW. In such an environment, users will be able to regard other remote users as inhabiting a (very large!) virtual hallway or office spanning several physical offices and organizations.

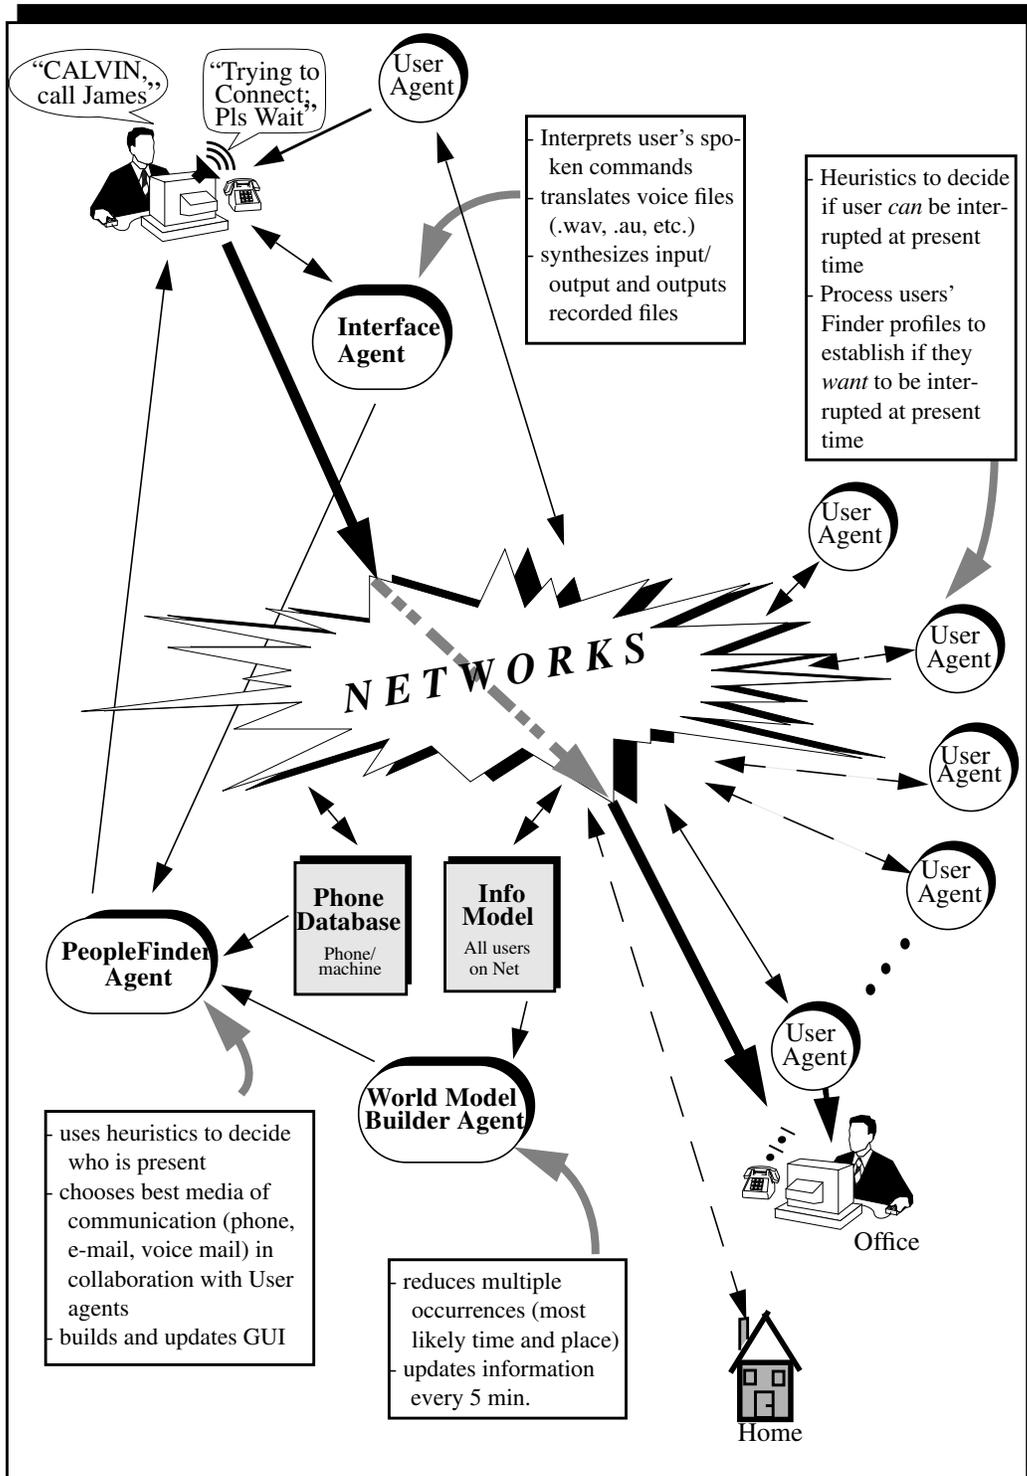


Figure 3. PeopleFinder: System Level View.

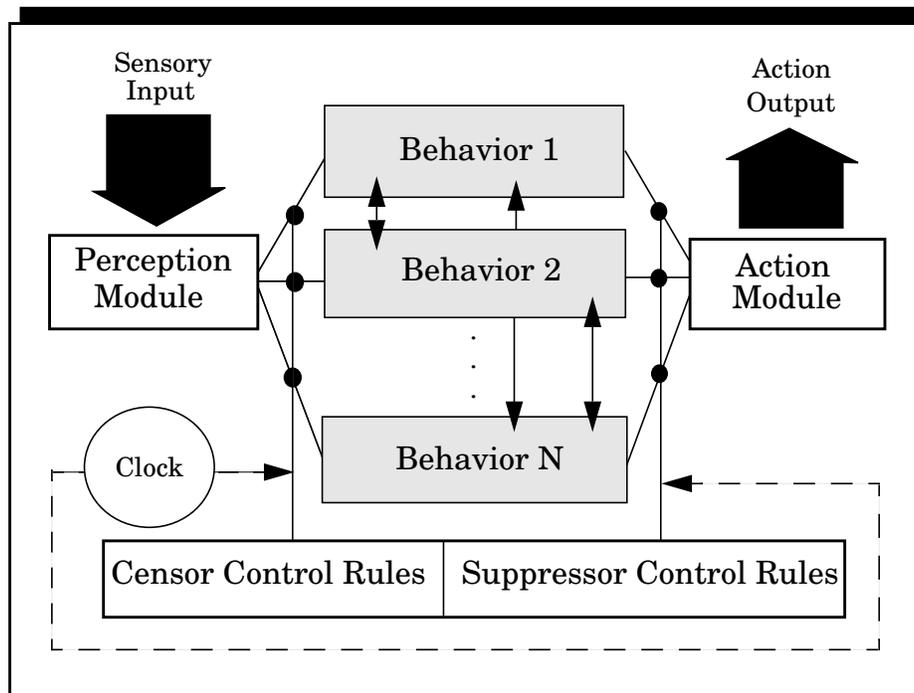


Figure 4. CALVIN: Control Framework.

resulting from shared access to the different presentation and communications resources used by the system (see Werkmán's BN negotiation-based conflict resolution work for related issues [Wer94]); adding privacy features so that individual users may control how much information regarding their activities and whereabouts may be accessed through the PeopleFinder; and formalizing the various rules used by the PeopleFinder to combine multiple media with multiple modalities for both human-computer interaction and user-user communication (much along the lines of the work of Arens *et al.* [AHV93] on allocating multiple media).

References

- [AHV93] Arens, Yigal; Hovy, Eduard; and Vossers, Mira. On the Knowledge Underlying Multimedia Presentations. In Mark T. Maybury, editor, *Intelligent Multimedia Interfaces*. AAAI Press: Menlo Park, CA, 1993.
- [Bux94] Buxton, William. The Three Mirrors of Interaction: A Holistic Approach to User Interfaces. In L.W. MacDonald and J. Vince, editors, *Interacting with Virtual Environments*. New York: Wiley, 1994.
- [CCWB94] Cohen, Philip R.; Cheyer, Adam; Wang, Michelle; and Baeg, Soon Cheol. An Open Agent Architecture. In Working Notes of the AAAI-94 Spring Symposium on Software Agents, Palo Alto, CA, pp. 1–8, 1994.
- [FJ94] Faure, Claudie and Julia, Luc. An Agent-Based Architecture for a Multimodal Interface. In Working Notes of the AAAI-94 Spring Symposium on Intelligent Multi-Media Multi-Modal Systems, Palo Alto, CA, pp. 82–86, 1994.
- [Fer92] Ferguson, Innes A. TouringMachines: An Architecture for Dynamic, Rational, Mobile Agents. Ph.D. diss., Computer Laboratory, University of Cambridge, Cambridge UK, 1992.
- [Fer95] Ferguson, Innes A. Integrating Models and Behaviors in Autonomous Agents: Some Lessons Learned on Action Control. In Working Notes of the AAAI-95 Spring Symposium on Lessons Learned from Implemented Software Architectures for Physical Agents, Palo Alto, CA, March 27–29, 1995.
- [FMF92] Finin, Tim; McKay, Don; and Fritzson, Rich. An Overview of KQML: A Knowledge Query and Manipulation Language. Available through the Stanford University Computer Science Department, Palo Alto, CA, March, 1992.
- [Luc93] Giachino, Luca. Activity Sensing Through Portholes Images: A Bridge between Passive Awareness and Active Awareness,

Technical Report, OTP-93-08, Ontario Telepresence Project, University of Toronto, Toronto, ON, August, 1993.

- [WHFG92] Want, R.; Hopper, A.; Falcao, V.; and Gibbons, J. The Active Badge Location System, ACM Transactions on Information Systems, 10(1):91–102, 1992.
- [Wei93] Weiser, M. Some Computer Science Issues in Ubiquitous Computing, Communications of the ACM, 36(7):75–83, 1993.
- [Wer94] Werkman, Keith J. A DAI Architecture for Coordinating Multimedia Applications. In Working Notes of the AAAI-94 Spring Symposium on Intelligent Multi-Media Multi-Modal Systems, Palo Alto, CA, pp. 93–97, 1994.

Acknowledgments

Production of this video was made possible by Daniel Gama-che and Paul Amirault, both of the National Research Council's Institute for Information Technology. We would also like to thank our various colleagues for agreeing to appear in the video.