Infrastructure for Multi-User, Multi-Spatial Collaborative Environments

Stephen Peters, Luke Weisman & Krzysztof Gajos

Artificial Intelligence Laboratory
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

http://www.ai.mit.edu

The Problem: Current research into intelligent environments focuses on a single user interacting with a single space. While this is a useful interaction, it ignores the propensity of people to work on problems in teams, as well as the utility of moving projects from workspace to workspace. In this project, we extend an existing agent-based environment to seamlessly handle multiple users and spaces.

Motivation and Previous Work: Recent research into intelligent environments [7, 3] has focused primarily on systems where a single user interacts with a single intelligent space, through the use of speech, gestures, or more mundane keyboard input. These systems often ignore handling multiple people, yet much interaction occurs in a conference or meeting room, with multiple people interacting with each other. In addition, these research efforts usually limit themselves to a single intelligent environment, ignoring the need for users to have their information and applications transferred with them as they move.

Attempts to achieve these capabilities without representing individuals as individuals per se, such as the KidsRoom project [1], creates systems that can only recognize users as aggregate concepts, and are limited to very simple feedback mechanisms, without the ability to target the feedback to one user. The EasyLiving project [2] incorporates the notion of individuals into its system, and can track multiple independent people within its room. However, it is still tied closely to one specific space.

Approach: This project extends the existing Metaglue [4] agent system, augmenting its capabilities to handle multiple users and spaces using a simple knowledge representation of users and spaces, and coupling it with an information base describing the capabilities of the spaces and roles for the users.

The first step is to split the set of agents currently running in the system into groups (referred to as "societies" in the Metaglue system), based on the entity (the user or space) on whose behalf they are operating. For each of these agent societies, there is an "ambassador" agent, which contains information about the represented entity, and allows the entity to exert some control over the agent cloud. This division is not quite as clear-cut as it might seem at first glance; imagine a user interacting with a web browser on a wall screen. Is the web browser operating on behalf of the user, or is it operating on behalf of the room? The answer bears heavily on the behavior exhibited when another person enters and tries to view a web page – if the existing browser is working on behalf of the user, a new window should be opened, otherwise the existing page should be shared.

Once agents are divided into their societies, we encode information about the real world entities into the ambassadors, so that the agents can utilize the information during operation. This can include configuration information requested by the user or space, resource management subsystems, and connections between the different societies. For example, even if the web browser agent above is operating on behalf of the room, it might still need to keep track of which user requested its current data. To accomplish this, the ambassador agent for the room can track the current users and the connections between them.

The result of this is that the details of a society become abstracted away behind the ambassadors, and agents within a society can be seen merely as the capabilities that the society presents. This makes it far easier for a user moving from one environment to another to transfer the display functions of his or her work (web pages being browsed, email agents) to another space, without being bogged down in the details of the environments. Indeed, most of the details can be encapsulated into a resource management framework [6].

Impact: Performing this kind of separation makes it easier to represent some more abstract concepts easily. For one, creating a representation of mobile devices into an existing space is actually simpler. They can simply exist as a virtual environment of their own, coordinating through a user with the environment at large. We call these "mobile spaces", preserving the abstraction between users and the devices that they are attempting to use, and making it possible to decouple the mobile devices from the user as well as the spaces they inhabit. Users can carry the devices with them, leave them behind, use them as extra displays or input devices, in as flexible a fashion as the resource management system allows.

In addition, it becomes easier to manage *ad hoc* groups of users for collaborative applications [5]. Temporary societies can be automatically created by a user's agents for the sole purpose of managing the collaboration. The ambassador for these "group societies" manages the interactions, and automatically destroys the societal information when the last client disconnects.

Future Work: As this is an ongoing project, some issues involved in the ambassador infrastructure are still being resolved. Initial tests integrating simplified ambassadors and simple resource management systems show promise, in that it does enable users to move state easily between environments. Work is currently proceeding on handling group societies and nested spaces.

Research Support: This work is funded by DARPA under contract number F33615-00-C-1702, administered by AFRL / IFSC, and by the MIT Project Oxygen.

References:

[1] Aaron Bobick, Stephen Intille, Jim Davis, Freedom Baird, Claudio Pinhanez, Lee Campbell, Yuri Ivanov, Arjan Schütte, and Andrew Wilson. Design decisions for interactive environments: Evaluating the kidsroom. In Coen [3], pages 7–16.

[2] Barry Brumitt, Brian Meyers, John Krumm, Amanda Kern, and Steven Shafer. Easyliving: Technologies for intelligent environments. In *Handheld and Ubiquitous Computing: Second International Syposium*, pages 12–29, 2000.

[3] Michael Coen, editor. *Proceedings of the 1998 AAAI Spring Symposium on Intelligent Environments.*, Menlo Park, California, 1998. American Association for Artificial Intelligence, AAAI Press.

[4] Michael Coen, Brenton Phillips, Nimrod Warshawsky, Luke Weisman, Stephen Peters, Krzysztof Gajos, and Peter Finin. Meeting the computational needs of intelligent environments: The metaglue system. In Nixon et al. [7], pages 201–212.

[5] Clarence Ellis and Jacques Wainer. *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*, chapter Groupware and Computer Supported Cooperative Work, pages 426–452. MIT Press, Cambridge, MA, USA, 1999.

[6] Krzysztof Gajos. Rascal - a resource manager for multi agent systems in smart spaces. In *Proceedings of CEEMAS 2001*. Springer, 2001. To appear.

[7] Paddy Nixon, Gerard Lacey, and Simon Dobson, editors. *Managing Interactions in Smart Environments: 1st International Workshop on Managing Interactions in Smart Environments (MANSE '99)*. Springer-Verlag, London, 1999.