DESIGNING AN ECOLOGY OF DISTRIBUTED AGENTS

EXTENDED ABSTRACT FOR MASTER'S THESIS PROPOSAL

December 1, 1997

Nelson Minar

<nelson@media.mit.edu>

http://www.media.mit.edu/~nelson/

Statement of Problem

The potential for the Internet as an environment for computation is only beginning to be realized. New software design approaches are needed to create flexible distributed computer systems. This thesis will explore possibilities for Internet applications by creating an *ecology of distributed agents*, a new paradigm for designing Internet software systems.

The representation of people's presence in the Internet is an interesting and important distributed application. Because data about users' activities comes from many different computers, visualizing presence requires a distributed system. Diverse representation approaches are possible, variations in both what data is collected and how it is displayed. Therefore, an effective presence application must be flexible, open-ended, and configurable. The problem of representing presence will benefit from novel approaches for building distributed applications.

Proposed Solution

I will build a system that represents peoples' presence by creating an ecology of agents. The application will consist of a population of different types of agents distributed on participants' machines. Agents will both manage and represent data about users. One genus of agents will collect data about users' activities, another genus will make graphical representations of that data, and a third will act as messengers to help agents find each other and communicate. Each genus will have several species of agents that carry out their roles in different ways: alternate data sources or visualizations, for example. Individual agents will be relatively simple; their complex interactions will implement the distributed application.

The system is designed as an ecology of agents. Individual agents are analogous to organisms in an ecosystem. The agents' environment is the network of servers running on participants' computers. Servers provide resources in the form of CPU cycles, memory, and data about user activity. Users configure their personal environment by supporting or killing agents running on their own server. The ecosystem as a whole maintains a diversity of agents while still allowing users to control their own systems.

Expected Results

Building an ecology of distributed agents requires solutions for many interesting technical problems. Distributed object facilities for Java are available that provide a basic infrastructure for agents. These existing software libraries will be augmented with mechanisms to allow agents to find each other, discover each other's capabilities, and coordinate their interactions. The design of the agent environment will require specifying the resources a server provides, creating a coherent network of multiple servers, and designing tools for users to manage agent populations.

The primary result from this research will be the working application itself: an adaptive, distributed system for representing people's presence in the Internet. After the software is written and deployed to a small group of users, the behavior of the system will be studied to understand how the ecology metaphor works in practice. The ultimate goal of this research is to explore and evaluate a new way of designing Internet software. A presentation of the implementation of the system will show the elegance of the distributed agents approach. An analysis of the behavior of the system will report on the effectiveness of building an ecology of distributed agents.