Mobile Agent-Based Search for Service Discovery on Dynamic Peer-to-Peer Networks

Evan A. Sultanik*
eas28@cs.drexel.edu

Dynamic multi-agent planning has many shortcomings in real-world domains. Specifically, it is not clear how services can be discovered and located in dynamic non-fault-tolerant networks. Examples of such include multi-hop ad hoc wireless and dynamic peer-to-peer networks. The nature of ad hoc and peer-to-peer networks inherently imposes the restriction of communication to one's "neighbors." Therefore, communication over multiple hops in the network is difficult without facilities such as ad hoc routing. In these networks, the problem of locating a service is compounded when services themselves can be mobile.

In some domains mobile agents can be service providers. For example, consider a group of helicopters deployed on a battlefield, as in the synthetic aircraft domain (Tambe 1998). A certificate authority might be required for secure communication between helicopters; this server could be encapsulated by a mobile agent capable of reasoning about the network (Artz, Peysakhov, & Regli 2003). The agent might then continuously migrate to portions of the network with low volatility. For instance, the agent might migrate to helicopters less likely to be removed from the network. To improve performance and minimize latency, heuristics for such migration might include proximity to the geographic center of the group or association with the centroid of the network topology graph. Therefore a method for pro-actively tracking the location of services in dynamic networks is required.

No fixed memory deterministic algorithm can locate a service in a network in a fixed amount of time (Kirousis *et al.* 2000). We propose a fixed-memory randomized method for approximating the location of a service in a dynamic network with a probabilistic certainty in a fixed amount of time.

Deploying a set of service monitoring agents to randomly walk the hosts of a dynamic peer-to-peer network can provide an accurate means of service discovery. The agents' task environment, a dynamic peer-to-peer network, is stochastic, dynamic, and continuous; there exists a delay between actual topology changes and the propagation of knowledge of these changes throughout the network. The agents do not have a goal, per se; their sole purpose is to randomly walk the network gathering information. Agents' percepts are comprised solely of the set of services available at the current host and the set of hosts neighboring the current host. Agents' actions are comprised solely of hopping to a neighbor host from their current host. At each host agents query for services, storing these data in memory (along with a timestamp). The agents' itineraries are dictated by the network; successor hosts for migration are selected randomly from the set of available neighbor hosts in the network. The agents' itineraries are dictated by the network; successor hosts for migration are selected randomly from the set of available neighbor hosts in the network.

Each host on the network is assumed to have a local state description represented by a tuple containing the following elements:

ν - probability an agent will be at host h;
η - number of instances of a service s;
$|\mathcal{H}|$ - cardinality of the set of hosts;
$|\mathcal{A}|$ - cardinality of the set of agents;
ℓ - average time needed for an agent to hop from one neighbor to another; and

Many of these elements need neither be calculated nor known a priori; they can be inferred from or provided by the service monitoring agents.

This approach has three important advantages over other alternatives, such as naïve message passing and broadcast:

1. a minimum of network bandwidth is used, and bandwidth usage scales linearly; this is an important issue for resource-constrained mobile devices and large-scale peer-to-peer networks;
2. services need not register themselves; and
3. properties of random walks are relatively easy to mathematically model and likewise make inferences upon.

Intuitively, the less frequently a host is visited by service discovery agents that have seen a service in some period of time, the less likely that service is currently available. Exploiting properties of random walks, such as limiting distributions, we have developed a mathematical formulation of agent movement in the network and an algorithmic technique for approximating the parameters of the formulation. Using this model, we can predict the availability of services with a high level of accuracy. This allows for time-critical reasoning and probabilistic inferences to be made upon the agent system.

Our current focus is on using our method for propagating service location information throughout the network as a

*Advised by Dr. William Regli

heuristic for mobile agent-based search. For example, each host could cache the data brought to them by the random walking service discovery agents. In doing so, each host would develop an index (or "belief") of the locations of services. These beliefs will become more accurate in conjunction with a host's proximity to the service. Therefore these beliefs, along with the timestamp of when the agent last saw the service, can be used as a hill climbing search heuristic when tracking down the service.

Our current research is focusing on the feasibility, admissibility, and accuracy of this heuristic, and its integration into the Secure Wireless Agent Testbed (SWAT) (Sultanik *et al.* 2003). This includes live tests with a dozen resource-constrained mobile computing devices, communicating over an ad hoc wireless network.

References

Artz, D.; Peysakhov, M.; and Regli, W. 2003. Network meta-reasoning for information assurance in mobile agent systems. In *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence*, 1455–57.

Kirousis, L. M.; Kranakis, E.; Krizanc, D.; and Stamatiou, Y. C. 2000. Locating information with uncertainty in fully interconnected networks. In *International Symposium on Distributed Computing*, 283–296.

Sultanik, E.; Artz, D.; Anderson, G.; Kam, M.; Regli, W.; Peysakhov, M.; Sevy, J.; Belov, N.; Morizio, N.; and Mroczkowski, A. 2003. Secure mobile agents on ad hoc wireless networks. In *The Fifteenth Innovative Applications of Artificial Intelligence Conference*. American Association for Artificial Intelligence.

Tambe, M. 1998. Implementing agent teams in dynamic multi-agent environments. *Applied Artificial Intelligence* 12(2–3):189–210.