



EXPLOITING EXPERTISE THROUGH KNOWLEDGE NETWORKS

Michael N. Huhns • University of South Carolina • huhns@sc.edu
Larry M. Stephens • University of South Carolina • stephens@sc.edu

If I want to find out how to increase the bandwidth into my office, I ask my network administrator, but if I want a new chair, I don't know whom to ask. My secretary, however, will know, so I ask her. In this way we all depend on each other, because few people know all they need to know to get their jobs done. Studies indicate that people who are more successful have faster networks of more capable experts, and they access this expertise in one-on-one interactions.¹ So *whom* you know is still as important as *what* you know.

In today's large or extended enterprises, where frequent personnel changes make it difficult to conduct business in such a direct way, success requires an ability to exploit the cumulative knowledge of a widely distributed and diverse workforce.² Moreover, the basic problem of knowing whom to ask for help has another side: you too have expertise in some area, and you need to find the people who can benefit from it.

Many companies are trying to facilitate such connections with some combination of Lotus Notes and an intranet portal navigated by a search

engine. If each employee has a home page on the intranet that clearly spells out that person's responsibilities and capabilities, then a search engine can compute an index for the home pages and store it in an online database. For example, along with 7,000 internal Web sites, Hewlett-Packard's intranet has a database of company-wide expertise. A properly conducted search can match the right person to the right task. In some companies, Lotus Notes is then used to keep track of incomplete tasks, and properly motivated employees can browse the list of tasks to find some they can work on.

Unfortunately, in practice, combining an intranet portal and Lotus notes is insufficient, because it does not support the way most of us work. We don't randomly browse directories to find someone who might be able to help us; instead, we're accustomed to finding the right people through our personal offline networks. Such networks of expertise tend to have very few levels, and they are not fully connected. In fact, to be useful, they need paths that are short and fast. As Singh pointed out,

there are only a few degrees of separation between most people in social networks,³ and we can assume this pertains to business as well. One-level referrals are the basis for some very successful search engines, such as Direct Hit⁴ (<http://www.directhit.com>).

Agent-Based Knowledge Network

In the long run we need a more comprehensive solution. The necessary capabilities are

- categorizing (the ability to classify Web pages and other unstructured data automatically),
- hyperlinking (the ability to add to each item of information appropriate pointers to other relevant items of information),
- alerting (the automatic notification of users and agents to new information that might be of interest to them), and
- profiling (the construction of models of users and agents to describe their interests and expertise).

This last capability is the most important for a knowledge network because it involves integrating statements of work, contracts, plans, and corporate strategies with structured data to characterize an enterprise's objectives and work.

The system architecture must include

- brokers that manage the metadata relating applications, agents, systems, and people,
- search engines,
- ontology servers to reconcile the semantics of the different components that make up the intranet,
- knowledge bases for each active participant in the system, and
- agents (of course!) to provide the proactive behavior needed to make the knowledge network an active collaborative service.

How close are we to achieving such a solution? It appears that all of the individual pieces are available, but they are just not integrated with agents into a complete system. Here is the current state of portals, groupware, and corporate knowledge management.

Intranet Portals

Corporate intranets are often controlled by many independent groups. As a result, intranet Web sites are built according to different standards, use different vocabularies and taxonomies, provide incompatible interfacing techniques, and are not coordinated. The information is accessible, but it is not comprehensible or consistent. A corporate portal is an attempt to improve on this situation by providing a single coherent point of access to information.

Several different kinds of portals are possible. The simplest is an information portal that allows access to Web pages, aided by a conventional search engine. A second kind enables users to collaborate by establishing virtual project communities with conferencing, workflow, query tools, and document management. A third type, the kind suggested here, enables expertise to be networked throughout an enterprise. Such a portal can be active, pushing knowledge to where it is needed and where it can be exploited. An early version is available from Intraspect Software Inc (<http://www.intraspect.com/index.htm>).

Communities of Practice

Groupware and portal solutions tend to be one-to-many. What is needed instead is a network of computational agents and people that can keep track of each other and use each other's strengths.⁵ But software such as MS Exchange or Lotus Notes doesn't support agents, and agent technology doesn't support people. Agents communicate via messages in languages such as KQML or FIPA, and people cannot be expected to communicate with agents or each other in such languages.

A better approach in many companies has been to set up "communities of practice," allowing workers to share knowledge on a particular topic.⁶ The necessary investment of time and leadership to establish and maintain such a community can pay off by fostering person-to-person interaction. AmberWeb at Xerox is one such community-of-practice Web site. Currently, 30,000 researchers, planners, and marketers use it to exchange information about the latest corporate research.

System of the Bimonth



ExpertCity.com (<http://expertcity.com>), scheduled to launch this fall, is the perfect site for those who keep giving away what they should rightfully be getting paid for. It is particularly appropriate for most engineers and computer scientists, who so enjoy the challenge of solving a problem that they forget to charge for it.¹ It also promises to be a useful meeting site for consultants and prospects.

Here's how it works. First you register as an expert in a given technical domain. Then when neophytes come to the site to ask a question, you and other experts can bid to answer their question for a fee. You can make only one bid per question, but if selected, you have a contract with the questioner. The resultant interaction is like a contract net, one of the first interaction protocols developed for multiagent systems. With the holidays coming up, this is a nice way to earn some extra money on what essentially is a virtual help desk driven by an auction of your services. Check it out!

Reference

1. M. Lattig, "Getting Business Smarts," *Infoworld*, 2 Aug. 1999, pp. 32-33.

Xerox's other version of this concept is Eureka, a system of more than 5,000 service tips that aid 22,000 technicians worldwide. Eureka started from small, informal gatherings of repair technicians who shared stories and insights about how they had solved customers' problems. The insights were captured in a knowledge base and then expanded into a system that now supports a global community of technicians. To support smaller groups needing to collaborate, Xerox provides DocuShare, which enables the creation of virtual workspaces of shared information for pre-identified group members.

Knowledge Management

AT&T enhances its intranet portals through knowledge management applications. One such application enables employees to analyze the status, returns, and future payouts on their pension investments via "what-if" analyses, while another aids meeting planning. These applications are individual and not collaborative, however. What's missing is an agent-based application for expertise.

Bottom Line

According to International Data Corp. in Framingham, Massachusetts, Fortune 500 companies lose an average of \$12 billion annually because they cannot exploit all the internal and external information resources

available to them. Fixing this will require knowledge networks comprising documents and document management systems, data and data warehouses, enterprise information portals, search engines, decision-support tools, collaboration tools, and intelligent software agents. Is the investment worth it? What's *your* company's share of \$12 billion?

REFERENCES

1. R.E. Kelley, "How to be a Star Engineer," *IEEE Spectrum*, Oct. 1999, pp. 51-58.
2. E. Borin, "Who You Know—and What They Know," *Knowledge Management*, Sept. 1999, pp. 85-86.
3. M.P. Singh, "Degrees of Separation," *IEEE Internet Computing*, July-Aug. 1999, pp. 4-5.
4. J. Tanaka, "The Perfect Search," *Newsweek*, 27 Sept. 1999.
5. S. Boutwell, "Share the Knowledge," *Information Week*, 13 Sept. 1999, pp. 65-70.
6. D. Deckmyn, "Human Interaction Key to Knowledge," *Computerworld*, Oct. 4, 1999, p. 55.

Michael N. Huhns is a professor of electrical and computer engineering at the University of South Carolina, where he also directs the Center for Information Technology.

Larry M. Stephens is a professor of electrical and computer engineering at the University of South Carolina. His current research interests are multiagent systems and ontologies.