

# COMMUNITY-BASED SERVICE LOCATION

HOW VIRTUAL COMMUNITIES AND SOFTWARE AGENTS CAN BE USED TO PROVIDE AND LOCATE TRUSTWORTHY SERVICES ON THE INTERNET.

Although both computing and communications capabilities are improving exponentially, there has been a more rapid increase in available bandwidth than in computing power over recent years [6]. Naturally enough, while the network-related improvements create new possibilities for applications, the improvements and especially their differential, render obsolete traditional ways of thinking about telecommunications and computing.

There is an ongoing debate occurring among two main factions within the telecommunications industry [1, 4]. On one side are traditional telephone companies, who advocate intelligent networks with increased intelligence embedded within the network and controlled by the telecommunications providers. The intelligence would be reflected in features such as for caller identification or call forwarding. On the other side are the proponents of (what may be called) stupid networks, who take their inspiration from the Internet. In this view, as bandwidth becomes plentiful, intelligence will propagate to the edges of the network and the network itself will provide no more than bit transport. The overwhelming advantage of stupid networks is that, like the Internet, they naturally support heterogeneity

and extensibility. End users can choose whichever applications they like and invoke whichever services they prefer without requiring consistent changes throughout a large network.

Consider telephone directories. White pages and yellow pages are essential for making a telecommunications system practical, but are classical centralized functionalities. Internet portals are a close analogue of telephone directories, and provide a one-size-fits-all solution to users' information needs. However, just as advances in computing are driving manufacturing from mass production to

mass customization, advances in communication are driving information access from large portals toward personal contacts. Large portals won't be eliminated, but will increasingly be supplemented and superseded by personalized sources of information as people increasingly want to receive information and advice from those whom they know and trust.

Accordingly, in this article we consider the problem of service location. We study an approach that places the intelligence on the endpoints, enabling the users to locate desirable services based on trustworthy, personalized recommendations of their peers. The task is not only to locate a particular service, but also to locate a service that is rated highly by one's friends and associ-



MUNINDAR P. SINGH, BIN YU, AND  
MAHADEVAN VENKATRAMAN



ates, and their acquaintances.

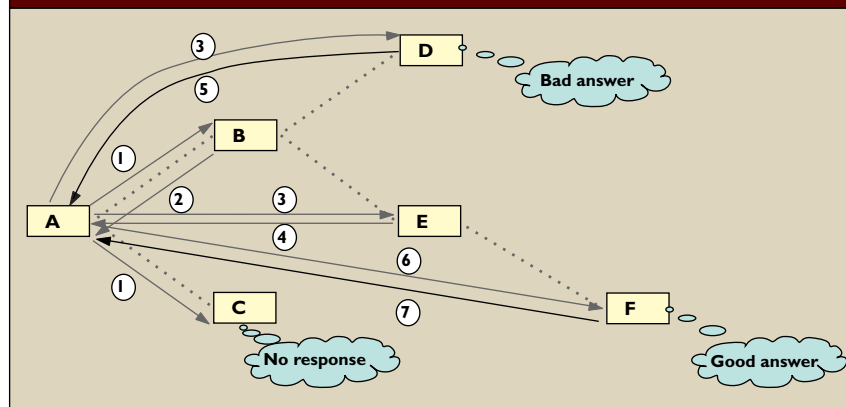
Open environments are characterized by having components that are autonomous (acting independently) and heterogeneous (designed independently), of dynamic membership (joining, changing, and leaving arbitrarily), and of large scale (numerous). These properties are most compatible with the stupid network doctrine described previously, and are best supported by an approach such as ours that is based on a peer-to-peer model of interaction among intelligent entities existing at the edge of the network.

### Communities and Agents

Let's review some key concepts and challenges. We define an online community as a set of interacting members or principals. The principals could be people, businesses, or other organizations. The members of a community provide services as well as referrals for services to each other. Our notion of services is general in that they need not be business services provided for a fee, but may be volunteer services, or not even services in the traditional sense, for example, just companionship or lively discussion about some topic. By the same token, quality of service includes not only the quality of the basic service but also of ancillary features, such as privacy. Thus a dentist who does a good job on your root canal, but tells everyone about it may be treated as offering a poor quality of service. In other words, the quality of service requested or provided is a multidimensional concept. The members of a community provide not only services but also referrals to each other. Referrals may be provided proactively or in response to requests.

Agents are persistent computations that can perceive, reason, act, and communicate. Agents are of interest here, because they represent different principals and mediate in their interactions. The agents assist their principals in evaluating the services and referrals provided by others, maintaining contact lists, and deciding or suggesting whom to contact for different services. In this manner, the agents assist their principals in finding the most helpful and reliable parties to deal with. The recommendations by the personal agents are based on a representation of how much the other available parties can be trusted. The agents build and manage these representations. To do so, the agents not only take into

**Figure 1. How the computation spreads as a query originates from an agent, referrals are sent back, and additional queries are sent out.**



account the previous experiences of their principals, but also communicate with other agents (belonging to other principals). Principals have full autonomy in deciding whether or how to interact with others. Because the agents participate on behalf of different principals, the agents appear as autonomous and heterogeneous.

The agents organize themselves into communities. The combination of personalization and user control over the flow of information leads to a community-based treatment of even staple telecommunications functionalities such as service location. Virtual communities build on real communities but make enhancements beyond the underlying real communities in terms of scale and precision. Manually maintaining extended communities would be a nuisance, but when agents can take over much of the grunt work, the community-based approaches are suddenly at a great advantage.

### Social Networks

The agents maintain a social network for their principals. In this network, the participants' naturally have reputations both for expertise (providing good service) and sociability (providing good referrals). Consider the service location scenario where agents help their principals find principals offering desired services. A query is a request for information about who provides a specified service. An answer can be a referral to another principal or even oneself, in which case there would be some additional interaction to actually provide the service. Figure 1 depicts the spread of computation between agents. As illustrated in Figure 1, agent A sends a query to its neighbors B and C; there is no response from C. Agent B sends back two referrals apiece to D and E, and E sends back a referral. The dotted lines in the figure indicate the relationship between agents, and the

numbered circles indicate the following:

- 1: Initial query by A to B and C (where C ignores);
- 2: Referrals to D and E given by B;
- 3: Query by A to D and E;
- 4: Referral to F by E;
- 5: Answer by D to A;
- 6: Query by A to F; and
- 7: Answer by F to A.

Queries originate from agents acting on behalf of their principals. Sometimes an agent may come up with a query proactively; at other times, the principal may explicitly indicate the need. An agent initiating a query decides the potential contacts to whom to send the query. Possibly after consultation with its principal, the agent sends the query to the agents for other likely principals. Agents filter incoming queries for their principals. In simple cases, the agent may itself be able to supply a referral on behalf of its principal. Even when prompted, the principal may autonomously either produce an answer or ignore the query.

In addition to or instead of just forwarding the query to its principal, the agent may respond with referrals to other principals; a referral if given is accompanied by a rating of the principal being referred. This is a valuable aspect of the social network, because it helps agents disseminate ratings of others and thereby to collectively construct the reputations of the members of the society. When the originating agent receives a referral, it decides whether to follow up on that referral. When the originating agent receives a service, it and its principal evaluate the quality received for future reference. This evaluation affects the originating agent's model of the expertise of the serving principal, and the originating agent's models of any agent who may have given a referral to the serving principal.

## Representations and Reasoning

The main questions for any work in computing are how the necessary information is represented and how it is reasoned with. We base our representations on the vector space model (VSM), a well-known information retrieval technique [7]. Queries, responses, and expertise are all modeled as term vectors in a multidimensional information space. The query vector is generated from the initial query. The expertise vector for each principal is estimated by the originating agent (based on previous interactions and referrals). We systematically compare a query vector with the expertise vectors of other principals to find the principal whose expertise is the most similar to the given query.

The VSM defines the similarity of two vectors as the cosine of the angle between them. Loosely follow-

ing this intuition, we define the similarity between a query vector and an expertise vector as the projection of the expertise vector on the query vector. The idea is that as far as a given query is concerned, only the expertise along the lines of that query is relevant. Further, the magnitude of the expertise matters—the more the better.

The relevance of a given principal to a given query depends not only on the similarity of the query to the principal's expertise, but also on what weight is assigned to the principal's sociability, which reflects how much we can trust the referrals produced by this principal. Accordingly, we define the relevance as a weighted sum of the expertise and sociability components. In order to make its key decisions, an agent must maintain the following information:

- To filter incoming queries, its principal's areas of expertise; and
- To decide whom to contact, models of peers, including their expertise vector and sociability; all neighbors and some non-neighbors may be modeled.

The VSM provides a basis for an agent to estimate the match with a peer's expertise vector. The models of peers—their expertise vectors and sociabilities—must, however, be learned. The models of peers are learned by the agent through the interactions of its principal with other principals. When a good quality of service is obtained, the expertise of the providing principal is revised upward as is the sociability of the principals who gave referrals to the good principal. When a poor quality of service is obtained, the revisions are downward. Ultimately, what every agent is trying to learn about others is their potential usefulness for finding services of the sort most often requested by its principal.

## Learning from Referrals

Let a principal  $p_i$  make a query. Suppose after a series of referrals, this query produces a service from principal  $p_j$ . The originating principal,  $p_i$ , will evaluate the service obtained. The evaluation will be used to upgrade or downgrade the expertise of the principal who provided the service and the sociability of every principal who gave a referral that eventually led to this service. The extent of the upgrade and the downgrade is tuned to give the most credit to principals who provide good service or give referrals yielding short referral chains leading to good service providers:

- A good service. The estimated expertise for  $p_j$  is upgraded. For all intermediate principals whose referral led to this service, the sociability is

upgraded.

- A bad service. The estimated expertise for  $p_j$  is downgraded. For all intermediate principals, the sociability is downgraded.
- No response. This is treated the same as bad service with the estimated expertise set to a low value. That is, for any referral chain which peters out into nothing, all members of the chain are penalized.

All referral chains are considered equally. If  $p_j$  is referred to be many principals and produces a good quality of service, all referral chains gain in sociability. Conversely, if  $p_j$  produces a poor quality of service, all referral chains suffer in sociability (see the table appearing here). In general, the agents may keep track of peers other than just their neighbors. Periodically, they decide which peers to keep as peers, that is, which are worth remembering. The agents also decide which of their peers to promote to neighbors and which neighbors to demote based on the expected usefulness of all peers.

### Basic Results

Our approach is being implemented on a distributed system of desktop and mobile platforms. While practical deployment would provide some validation for our approach, it would limit the kinds of situations we can study. To better understand the principles underlying community-based service location, we have conducted several experiments on a simulation of the setup described previously. All the heuristics used in the simulation are as in the actual system. However, the queries and responses are generated automatically, rather than by humans. An agent's queries are generated by perturbing its interest vector, reflecting the intuition behind the interest vectors.

Our experiments involve between 20 and 60 agents with interest and expertise vectors of dimension five. The agents send queries, referrals, and responses to one another, all the while learning about each others' interest and expertise vectors. The number of neighbors per agent is limited—in our case to approximately four. The neighbors are the agents to whom a given agent may send a query or about whom the agent may issue a referral. The idea is that the agent should have few neighbors relative to the entire society of agents.

Some aggregate properties of a social network must be defined so we can study networks and their evolu-

tion experimentally. The effectiveness of a social network can be defined in terms of the likelihood of finding high service quality with the least number of messages. This leads us to define the quality of a social network as the sum of the similarity between each principal's interest and every other principal's expertise, but discounted by the branching factor and the length of the shortest path between them.

Learning by X about Y and Z when X asks Y; if Y refers to Z, then X asks Z.					
Situation		Y's exp	Y's soc	Z's exp	Z's soc
Y gives a good answer		↑	–	–	–
Y gives a bad answer		↓	–	–	–
Y refers to Z					
	Z gives a good answer	–	↑	↑	–
	Z gives a bad answer	–	↓	↓	–
	Z gives a good referral	–	↑	–	↑
	Z gives a bad referral	–	↓	–	↓

We obtained some useful results from previous experiments [9]. We describe these briefly and then go on to describe some more interesting results dealing with the structural properties of social networks:

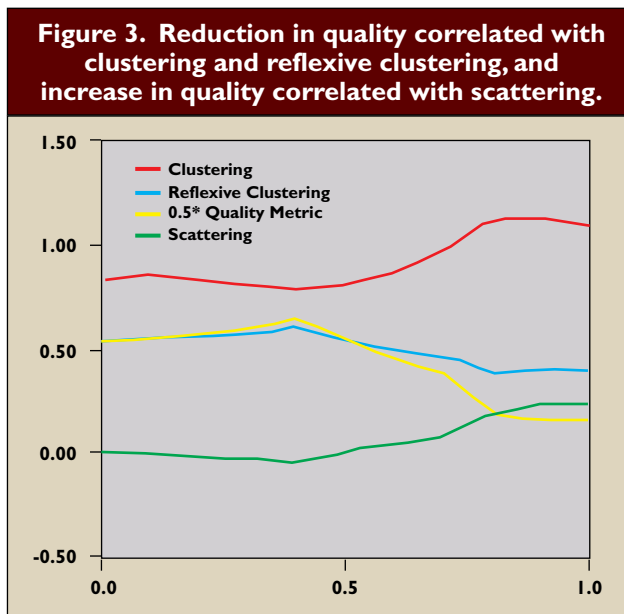
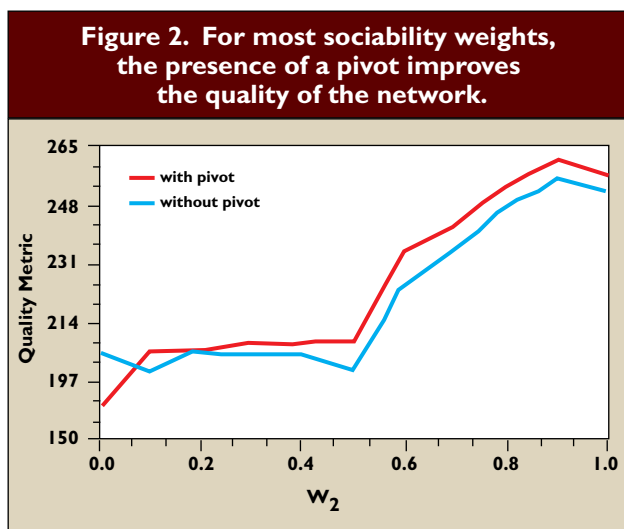
- The quality of the social network improves over time. The agents adaptively find better neighbors than they are randomly assigned at the outset.
- The social network stabilizes at an improved quality. The agents find neighbors with whom they can stay as long as their principals' interests don't change.
- Giving and taking referrals has a significant payoff. When referrals are given, the quality of the system is higher than otherwise.
- Giving consideration to others' sociability improves the quality of the social network, but an overemphasis on sociability (at the cost of expertise) can hurt.
- A new principal added to a stabilized social network will drift toward neighbors from which it obtains improved quality.

### Structural Properties

Clearly, not all the members of a social network are equally important in terms of their contribution toward the quality of the network. Some principals are more important in establishing contacts with the broader social realm. We call such principals "pivots." The agents in a community will tend to adapt to have neighbors whom they can trust. Thus subcommunities will tend to emerge. Importantly, pivots connect

across different subcommunities. The links from a principal to a pivot are termed weak ties and contrast with strong ties, which connect a principal to others in the principal's own community.

Sociologists have long known that weak ties are more powerful than strong ties, precisely because they go beyond one's own community [3]. Principals in the



same community would generally tend to have the same knowledge as everyone else in that community [2]. However, acquaintances from outside one's community will add greater value by having different knowledge and perspectives. Consequently, pivots are critical to a society.

We model pivots as agents that have a significantly higher out-degree (that is, number of neighbors) than average. Because of their higher out-degree, such agents are valuable to others and soon end up with a

high in-degree as well. Our simulations confirmed the hypothesis that the existence of a pivot agent significantly improves the quality of the social network as perceived by all agents in the network. Figure 2 illustrates this result—the existence of a pivot improves the quality at most weights of sociability.

A clique is a maximal complete subgraph of three or more vertices. Watts and Strogatz propose a clustering coefficient to estimate the cliquishness or the strength of clustering of a graph [8]. We define a directed graph version of this coefficient. The clustering of a vertex is the ratio of the actual edges among its neighbors to the possible edges among its neighbors. The higher this ratio the more well-defined the cluster around the given vertex. The clustering of a graph is simply the average clustering of its vertices.

Although a good approximation of social clustering, the preceding metric has two shortcomings. It rates a vertex as contributing to cliquishness merely if it has edges into a well-connected subgraph even if the members of that subgraph don't have edges back to it. Conversely, the metric also rates a vertex as contributing nothing to the cliquishness of the network if it has edges to vertices that don't know each other. We introduce another metric that avoids these shortcomings. The reflexive clustering of a vertex is the ratio of the actual to the possible edges involving that vertex or its neighbors. Now the given vertex itself is included in the counts of edges. The reflexive clustering of a graph is the average reflexive clusterings of its vertices.

Interestingly, although clustering and reflexive clustering have similar definitions, their difference is important. The difference reflects the existence of pivots in the social network, and leads to an interesting result. We define the cautious scattering of a graph as its reflexive clustering minus its clustering.

We found that both clustering and reflexive clustering have a similar relationship with the quality of a social network. If either kind of clustering goes up, the quality of the network goes down. This is illustrated in Figure 3—the quality of a network increases when its clustering and reflexive clustering decrease.

Interestingly, scattering has the opposite relationship with quality. This is also illustrated in Figure 3—the quality of a network increases when its cautious scattering increases.

Figure 3 shows the results for when interest and expertise were chosen independently. However, almost identical relationships are obtained even when interest and expertise are set equal or opposite to each other. Thus, our main claim regarding scattering and quality appears quite robust.

This is an important claim, because current approaches such as collaborative filtering are predicated on clustering users by similarity. When you have to predict a match based on anonymous averages, clustering would indeed be a safe bet, but when you can receive personalized recommendations, you would rather you receive them from people who are different from yourself. Clustering tends to increase the distance to useful experts because more links are used up within a small community.

## Discussion

In open environments, locating the right services is a major challenge. Our approach has some nice features. First, it is not intrusive on users and requires minimal actions and decisions by them. Second, it is not wasteful of resources—especially important for mobile applications. Third, it provides for users to naturally build a social network on top of the telecommunications infrastructure.

Our approach contrasts not only with traditional telecommunications approaches, but also with recent computing approaches such as collaborative filtering, employed at sites such as amazon.com. Collaborative filtering offers advice to a user based on the choices of users similar to the given user [5]. Collaborative filtering employs a client/server paradigm in which the server aggregates the choices of several clients (users) and decides what recommendations to offer to other users. This approach has the restriction of identifying the rating user to the server, while not revealing the source of recommendations to a user receiving recommendations. This is the inverse of what people really want: not to declare their ratings except to friends, but to know who provides ratings to them. Our approach, by contrast, is decentralized and lets the users control to whom they reveal their ratings.

We expect the greatest impact of our work to be for the problem of trust management. Trust is important wherever autonomous parties must interact. Importantly, trust is different from security. Security techniques such as passwords and digital certificates can assure you that a party you are dealing with is authenticated and authorized. However, security techniques do not ensure that the other party exercises its authorizations in a way that serves your interests. In other words, the security approaches simply place a low hurdle of legality that parties must cross in order to participate, whereas trust management makes them accountable even for the legal actions they perform.

While some aspects of security are best placed in the telecommunications network, trust can require such subtle combinations of features that it cannot be placed inside a centrally managed telecommunications

network. Trust must be handled from the edges of the network where different parties can build their reputations for trustworthiness in an application-specific or community-specific manner.

Peer-to-peer (P2P) systems, such as Napster and Gnutella, have started to emerge. These systems enable users to exchange information directly without going through a server. In their present form, P2P systems have narrow scopes, for example, the exchange of music files. (They are also currently controversial, because they can be used to violate copyright restrictions.)

Existing systems don't address the problems of finding high-quality, trustworthy peers. However, techniques for service location are critical to the expansion of P2P networks into newer domains. For example, how does one know that a copy of Lincoln's Gettysburg Address obtained from someone's computer is correct, and how does one know that an accountant is as good as he or she claims to be? Centralized service location would obviously not be appropriate. The best solution lies in the P2P networks themselves. Besides connecting people so they can exchange information and other services, P2P networks can also provide a substrate for community-based service location. It is in this form that we believe P2P networks will expand into broader domains of interest. **□**

## REFERENCES

1. Gilder, G. George Gilder on the bandwidth of plenty (as interviewed by Charles Petrie). *IEEE Internet Computing* 1, 1 (Jan. 1997), 9–18.
2. Gladwell, M. Six degrees of Lois Weisberg. *New Yorker* (January 11, 2000), 52–63.
3. Granovetter, M. The strength of weak ties. *American Journal of Sociology* 78, (1973), 1360–1380.
4. Isenberg, D.S. The dawn of the stupid network. *NetWorker* 2, 1 (1998), 24–31.
5. Resnick, R. and Varian, H.R. Recommender Systems. Guest editors' introduction to special issue of *Commun. ACM* 40, 3 (Mar. 1997), 56–58.
6. Robert, L.G. Packet switching or optical switching. *IEEE Internet Computing* 4, 1 (Jan. 2000), 50–51.
7. Salton, G. and McGill, M.J. *An Introduction to Modern Information Retrieval*. McGraw-Hill, NY, 1983.
8. Watts, D.J. and Strogatz, S.H. Collective dynamics of small-world networks. *Nature* 393 (1998), 440–442.
9. Yu, B., Venkatraman, M. and Singh, M.P. An adaptive social network for information access: Theoretical and experimental results. *Applied Artificial Intelligence*, 2001, To appear.

---

**MUNINDAR P. SINGH** (singh@ncsu.edu) is an associate professor at North Carolina State University.

**BIN YU** (byu@eos.ncsu.edu) is a doctoral candidate in computer science at North Carolina State University.

**MAHADEVAN VENKATRAMAN** (mvenkat@eos.ncsu.edu) is a software engineer at iScale, Inc., in Mountain View, CA.

---

This research was partially supported by the National Science Foundation under grant IIS-9624425 (Career Award).

---

© 2001 ACM 0002-0782/01/0400 \$5.00