Eight European Conference on Case-Based Reasoning, ECCBR06
September 4th-7th, 2006, Fethiye/Ölüdeniz, Turkey

The PwC Connection Machine: An Adaptive Expertise Provider

Mehmet H. Göker, Cynthia Thompson, Simo Arajärvi, Kevin Hua

PricewaterhouseCoopers LLP
Center for Advanced Research
Ten Almaden Blvd, Suite 1600
San Jose, CA 95113
`{mehmet.goker, cynthia.thompson, simo.arajarvi, kevin.k.hua}@us.pwc.com`

Abstract. The Connection Machine helps PricewaterhouseCoopers LLP (PwC) partners and staff to solve problems by connecting people to people. It allows information seekers to enter their question in free text, finds knowledgeable colleagues, forwards the question to them, obtains the answer and sends it back to the seeker. In the course of this interaction, the application unobtrusively learns and updates user profiles and thereby increases its routing accuracy. The Connection Machine combines features of expertise locators, adaptive case-based recommender systems and question answering applications. This document describes the core technology that supports the workflow, the user modeling and the retrieval technology of the Connection Machine.

1 The Power of Connected People

Information, knowledge and experience are key success factors and the most important competitive advantage for any business. However, most of this core corporate asset is in the heads of the employees and cannot be easily accessed, shared or distributed. Capturing and protecting it in documents (electronic or otherwise) is not only cumbersome, but the documents become rapidly outdated and the maintenance effort required to keep document collections up-to-date is formidable.

Furthermore, in the complex business scenarios of today's world, problem solving requires an increasingly large amount of specialized knowledge. It is nearly impossible for one individual to be an expert in every aspect of a company's business and deliver comprehensive solutions. Problem solving requires co-operation and the sharing of ideas and information. The size of a corporation and the collective knowledge of its employees are only valuable if these employees can share their information and cooperate. We believe that the best way to provide the most up-to-date and accurate information to those who seek it is by putting them directly in touch with the experts.

The PricewaterhouseCoopers Connection Machine is an application that enables employees to solve business problems by helping them obtain answers to their questions from knowledgeable colleagues. Rather than trying to extract information from experts and pointing information seekers to stale document directories, the Connection Machine matches incoming questions to the expertise profiles of users, routes questions to the experts with highest similarity, collects their answers and relays the answer back to the seekers. The application extends the personal network of employees to the entire firm and makes otherwise difficult to reach experts accessible.

2 Existing Approaches to Locating and Contacting Experts

2.1 Directory Systems

Most firms allow their employees to search for other colleagues by means of directories. Typically, these directories list the business unit, office phone numbers and addresses of employees, as well as some limited information about their background. Searches are usually performed by entering the (partial) name of the employee or by browsing to through the business unit structure of the firm. In terms of their functionality, these systems resemble phone books with a job categorization, similar to “yellow pages”.

If we know which employee we are looking for, directory systems are very useful for finding their contact information. However, most of these applications do not help to determine which employee might be knowledgeable on a *specific* topic [c.f. 1] and able or willing to answer our question. They also do not help to relay the question to the right person, to obtain an answer in a given timeframe, or to create a network of employees. Additionally, the data that goes beyond office location, department, phone numbers etc. is typically not centrally maintained and requires manual updates by the employees themselves. As such, the information is mostly outdated and its reliability rather limited.

Also, in personal interactions, if experts are not able to give an answer to a question, they typically refer the inquirer to another specialist from their personal network. A user looking for an expert in a directory system has only access to one level of experts and is at the mercy of the expert he/she contacts. People who have no representative profile in the directory system are beyond the reach of the seeker entirely.

Since standard directory applications do not provide the functionality to find an expert and ask a question easily, employees typically revert to the rather inefficient practice of sending emails to broad audiences in the hope of finding someone who is able and willing to help them.

2.2 Expertise Locator Systems

To answer the need for being able to access experts and ask questions, companies have developed so called Expertise Locator Systems (ELS). These systems try to find experts that are potentially able to answer a user’s question by matching the query to the expertise profiles of the employees [2, 3, 4, 5]. Some systems enhance the matching process by using the social connections between employees or collaborative filtering (e.g. [1, 6]). Depending on the application, they return a combination of potentially knowledgeable experts and related documents. It is up to the user looking for information to contact the experts and to get an answer to their question

Employees are normally represented by an expertise profile which, depending on the application, contains a limited number of structured attributes coming from an enterprise directory, a list of documents published by the employee, and general background information in free text or in a list of terms/noun phrases. The experts can update their profiles manually by adding new documents, modifying their background information and, potentially, the structured data. Responses to queries can be published and added to the profile as new documents as well. Some systems generate profiles automatically by analyzing emails and authored documents and extracting a set of terms. Users have to go through the terms to specify which ones represent areas that they would feel comfortable answering questions in.

Current expertise locator systems are designed to search for *people*. They match the user's question with documents and expert's profiles and display the list of matching experts to the users. The users, in turn, have to pick an expert from this list and contact them with the question. However, the goal of users who submit questions to an expertise locator system is not to find the name of colleagues but to find answers to their questions! The fact that a user has found the name of a potentially knowledgeable person does not mean that his/her question has been answered.

An additional weakness in current expertise locator systems is the lack of division between interest and expertise. Existing expertise locators analyze documents that have been authored by users and their emails to generate a profile to represent each user's expertise. If a user subscribes to an electronic newsletter out of interest in the specific topic, or writes a "Request for Proposals" (RFP) for vendors to respond to, he/she will be presumed an expert in that field.

3 Overview of the PwC Connection Machine

The Connection Machine extends the concepts of directory systems and expertise locators beyond the pure search for *people* and helps PwC partners and staff to get *answers* to their questions and to solve problems together. It leverages the personal networks and intelligence of PwC employees, facilitates collaborative problem solving, and fosters a work environment in which people are truly connected.

By answering questions rather than just locating people, the Connection Machine acts as a virtual, adaptive expertise provider. It combines features of expertise locators with adaptive case-based recommender systems and question answering applications.

Figure 1 provides a general overview of the interaction between the information Seeker, potential Providers and the Connection Machine.

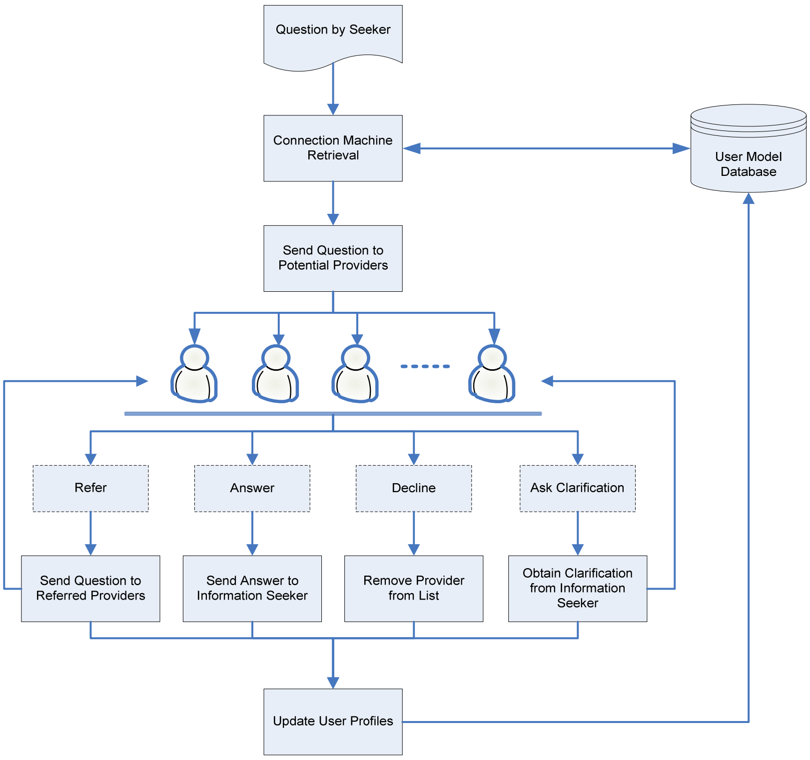

Figure 1: Overview of the application workflow in the PwC Connection Machine

An interaction with the Connection Machine starts with an Information Seeker entering a question in free text format, as if he/she were asking a colleague a question via email. The Seeker is also able to specify the urgency of the question, the name of a client the question relates to as well as additional, optional, structured information (e.g. knowledge domain, line of service, industry) to be used to locate appropriate potential Providers (Figure 2).

The Connection Machine processes the query, finds a set of suitable potential Providers and contacts them. The system only contacts potential Providers whose expertise levels for the given question are higher than the Seeker's and whose maximum number of questions per week has not been reached.

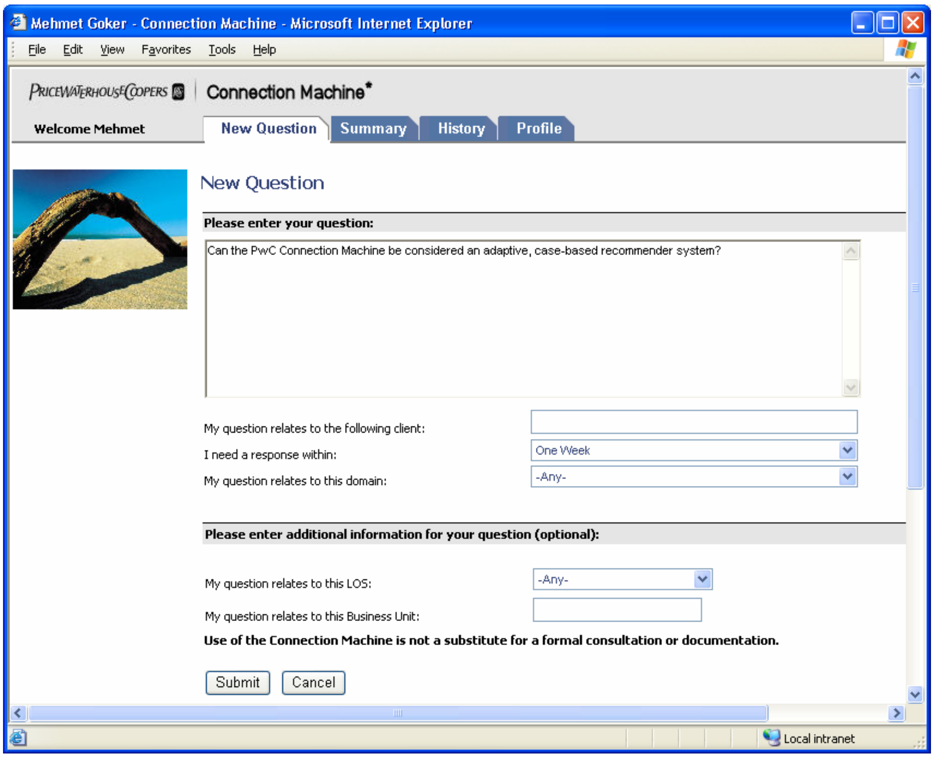

Figure 2: Web interface of the PwC Connection Machine

The number of potential Providers who will be contacted regarding a question is configurable. If the first set of potential Providers is not able to respond within the allocated time (a fraction of the time the Seeker needs the answer by), the system sends the question to a second batch of potential Providers. If no potential Providers can be identified or the Providers do not react, the question is sent to the Knowledge Administrator of the Domain for further processing.

Once the system identifies potential Providers, they are notified via email (Figure 3) and a visual indicator in the "Summary" page of the web interface, informing them that their expertise is needed. In addition to the question, the potential Providers are informed of the Seeker's contact information (e.g. name, line of service) and of the timeframe in which the question needs to be answered.

After receiving a question, the potential Providers may choose to respond either via web interface or via email. Potential Providers may offer an answer to the question; request additional information from the Seeker; refer the question to other potential Providers; or decline to answer. Once one of the potential Providers offers an answer or requests additional information, he/she becomes the "Provider" for the interaction. From this point on, the Connection Machine facilitates communicates between the information Seeker and the Provider and removes other potential Providers from the problem solving conversation by sending them email and removing the indicator in their "Summary" web page.

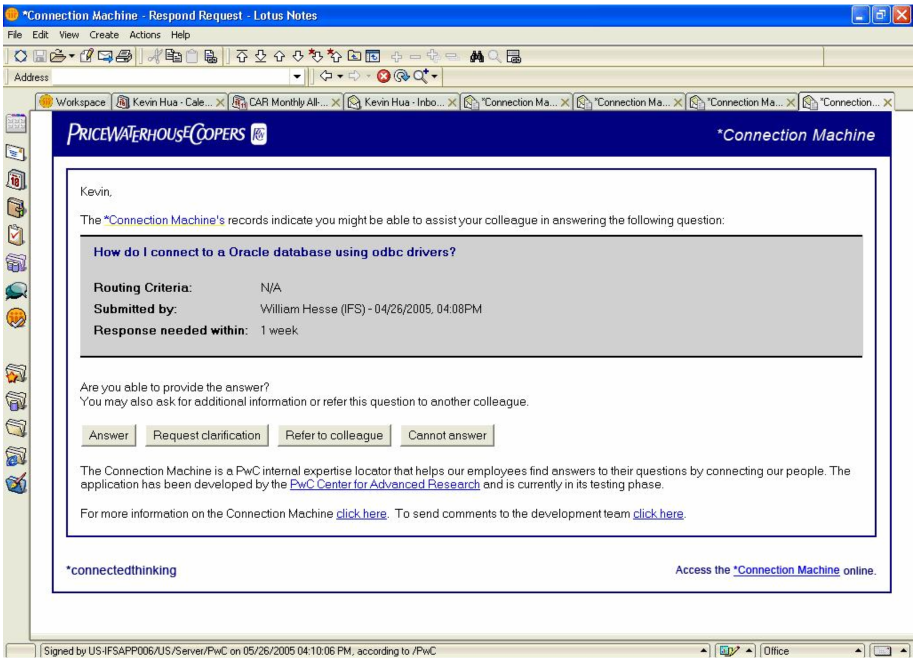

Figure 3: Sample email from the PwC Connection Machine

If a Provider chooses to answer the question, the Seeker is notified of the answer via email and a visual indicator in the web application (Figure 4). Upon receiving an answer, a Seeker can choose to accept it and close the request, ask a clarification question about the answer, or reject the answer and request a second opinion unless they have already done so.

The Provider can also ask for additional information that may be needed to answer the question. Once the Seeker provides additional information the Provider will be presented with the same options as when initially contacted by the Connection Machine (i.e. provide an answer, request question clarification, refer the question and decline to answer).

If a (potential) Provider decides that someone else from his/her personal network is better suited to answer the question, he/she may choose to refer the question. In this way the Connection Machine can learn about potential Providers who may have been missing from its initial set of profiles. The Seeker will not be made aware that the question was referred to another potential Provider as long as the initial Provider had not contacted the Seeker prior to referring the question (i.e. the provider did not request question clarification prior to referring the question).

The Provider can also indicate that he/she is not able to provide an answer to the question and specify the reason for declining to answer (e.g. “Too busy”, “Don’t know the answer”, “Independence conflict”).

If the question was declined by all contacted Providers, it will be sent to the Knowledge Administrator of the domain for further processing. The use of a Knowledge Administrator as a “backup” for answering or referring questions ensures that all questions entered in the Connection Machine are answered in a timely manner.

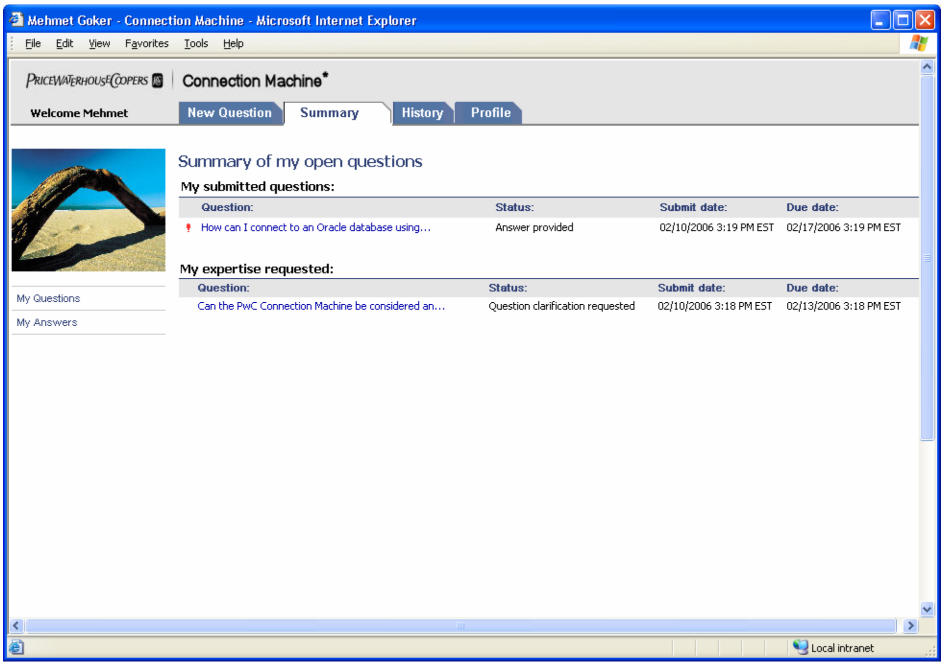

Figure 4: Open question summary page

4 Retrieval of Potential Providers in the Connection Machine

To execute the workflow described above the application needs to be able to determine who is potentially capable of answering the question by matching a user's query against information it has about other users. To achieve consistently high accuracy over a long period of time, the information of the users has to be updated with appropriate sections of the interaction on an ongoing basis. (Figure 5)

The technology we used to implement these functions in the Connection Machine is similar to User Adaptive, Case-Based Recommender Systems [7, 8, 9]. However, most recommender systems are geared towards selecting the best match out of a set of (mostly static) items and presenting it to the user. In the case of the Connection Machine, the items in the case-base are continuously evolving user models where each model contains multiple profiles. Rather than being the final goal, the retrieval process is an intermediate step and users, whose expertise profile matched the query, are utilized in the workflow to route questions. The resulting interaction between the Seeker and Provider is the desired outcome for the application.

The user modeling in the Connection Machine is not geared towards influencing the similarity metrics, the user interaction or the user interface of the application.

Neither can it influence the solutions a Provider may offer to a Seeker. The case-base of the Connection Machine is a collection of user models which are constantly maintained and updated by unobtrusively observing the user's interaction with the system [cf. 10] and, which are then used to select the users that will participate in the workflow as potential Providers.

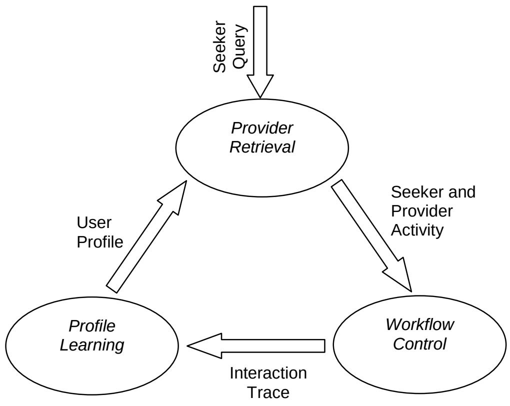

Figure 5: Profile maintenance and usage in the Connection Machine

5 The User Model of the Connection Machine

5.1 Interest, Expertise and Referral Profiles

A user model (i.e. case) in the Connection Machine contains three types of user profiles, each of which captures one aspect of a user's preferences or capabilities:

- The *interest profile* denotes the topics a user is interested in. It is updated from the questions a user asks and the associated clarifications.
- The *expertise profile* represents the topics that the user is knowledgeable in. This profile can be initialized from documents the user authored, a resume, prior engagement histories or similar. It is updated with the questions a user could answer, the answers he/she provided, and all associated clarification conversations.
- The *referral profiles* represent the topics in which the user is able to refer questions. It is updated with the questions, any clarification conversation associated with them, and any referral comments.

We also make a distinction between *positive* and *negative* profiles. Negative profiles contain information that the user does not want to be associated with. Users of the Connection Machine can choose to "Opt out" from a question by stating that they no longer wish to receive questions similar to the current one. This information is used to update the negative profiles. The use of negative profiles thus brings the total number of profile types available to six. In the following, we focus on "positive" profiles but all techniques apply to negative profiles as well.

Having six different profile types in a user model allows us to route questions with higher accuracy. It also enables us to generate reports on the distribution of interest, expertise and referral capabilities within the firm, to generate communities of interest or expertise and to route relevant documentation to inform group members or for review.

5.2 Data Sources for User Models

User profiles of directory systems and expertise locators are typically generated from structured data from sources such as enterprise databases, or from unstructured sources such as documents authored or read by the user. The resulting user profiles can be structured (e.g., a list of attributes and their associated values), unstructured (e.g., a list of terms, a collection of documents), or a mix of these representations.

While structured representations provide benefits in terms of retrieval accuracy and standardized vocabulary, they are very difficult to create and to maintain (both for the company and the individual). Since the Connection Machine operates in an environment where structured information is already generated and maintained for directory systems, we can utilize this information where available and concentrate our research on the generation of profile information from unstructured data sources and on the best way to represent it.

The information source we consider for the generation of user profiles is the set of documents (e.g., resumes, whitepapers, emails) that the user has authored, read or received. As such, the problem we are targeting is how to use *unstructured* data to generate and represent an expressive, flexible, and easy to create and maintain user profile.

5.3 Representation of Document or Term Based User Profiles

The representation that is used to store the content of the user models has a significant impact on the capabilities, flexibility, maintainability and learning abilities of the system.

In today's expertise locator systems, information from unstructured source documents is typically captured in profiles that are either based on term statistics extracted from discrete documents, or from term statistics based on the union of all documents associated with a user. In the first approach, the profiles are collections of term statistics per document and experts are ranked based on their respective total number of documents that are similar to the user's query. In the second approach, the profile is a set of term weights extracted from the union of documents for an expert and the similarity is computed between the query and the term weights for this collection.

We observed that the focus on discrete documents and their terms neglects the fact that the documents associated with a person can represent different facets of a bigger picture. A person who discusses the topic of "Hybrid Engine Performance" in one document and “Engine Emissions” in another is highly likely to be knowledgeable on “Hybrid Engine Emissions” or potentially “Engine Performance and Emissions” as well.

We also found that basing the profiles purely on terms (from discrete documents or the collection of documents) and neglecting the impact of phrases and relationships between terms can reduce accuracy and results in routing errors. For example, a person mentioning “Captive” at the beginning of a document, “Insurance” in the middle and “Bermuda” at the end may not be knowledgeable on the topic of “Captive Insurance arrangements in Bermuda”. On the other hand, a person who wrote a document with exactly this title and who did not frequently mention “Bermuda”, “Captive” and “Insurance” in the body of the document might be better able to help the Seeker than the travel agent who wants to sell an additional “insurance policy for a captivating trip to Bermuda”.

5.4 Lattice Based User Profiles

To address the problems described above, we devised a profile representation in which interest, expertise and referral profiles of users can each be incrementally learned and that captures a unified summary of the individual’s knowledge, crossing document boundaries. Using the documents associated with an individual we generate a profile for them that contains each term in these documents as well as the distance based co-occurrence weight between terms. Graphically, this represents a lattice[1] in which the nodes are the terms with their associated term weights and the links between the nodes are the co-occurrence weights between two terms. Below, we call any such group of linked terms a phrase.

Table 1: Notation used in the Lattice Representation.

For Terms	For Phrases	Description
N	N	Number of profiles of a particular type
$f_p(i)$	$f_p(i,j)$	Frequency (occurrence count) of term i or phrase i-j in profile p
$f_{max\ p}$	$f_{max\ p}$	Maximum term (or phrase) frequency in profile p
$pf(i)$	$pf(i,j)$	Profile frequency: number of profiles containing term i or phrase i-j
$ipf(i)$	$ipf(i,j)$	Inverse profile frequency of i th term or phrase i-j
$w_p(i)$	$w_p(i,j)$	weight of term i or phrase i-j in profile p

One lattice is built for each of the six profile types to represent a given user. We generate a lattice by starting with the available documents, appropriate for the profile type being built, that are associated with each user. Each document is divided into tokens, synonyms are processed, terms are stemmed, and stop words are removed. Once the documents have been pre-processed, we compute term weights, phrase distances, and other relevant statistics needed to create the profiles. Note that any document can be processed in this manner, including questions, answers, and other conversations between users.

[1] This is not a lattice in the mathematical sense.

Table 1 shows the notation used here to describe the process of computing the weights for each term and phrase. We begin with the terms. Each term in each user's profile is associated with a term weight. Our process for computing individual term weights is the same as standard TF-IDF (term frequency / inverse document frequency) approaches used in information retrieval (IR) [11], but is applied to profiles containing multiple documents rather than individual documents as in IR. Thus, the weight of term *i* in profile *p* is calculated as:

$$w_p(i) = tf_p(i) \times ipf(i) \tag{1}$$

In this equation, the normalized term (or phrase) frequency $tf_p(i)$ and the inverse profile frequency ipf_i are calculated as follows:

$$tf_p(i) = \frac{f_p(i)}{f_{\max p}} \tag{2}$$

$$ipf(i) = \log(1 + N/pf(i)) \tag{3}$$

To generate links between the terms (the "phrase" weights) we gather all pairs[2] of terms that occur together in a sentence, usually using a window size to restrict the number of pairs considered. Each link is then connected by a weight proportional both to the number of times the terms occur together in the profile and the number of intervening words between the terms. Thus, in the formulas above $tf_p(i)$ becomes $tf_p(i,j)$ and $ipf(i)$ becomes $ipf(i,j)$ with all formulas being adapted according to Table 1. For the case of phrases, we compute the frequency $f_p(i,j)$ by using a distance-based frequency count:

$$f_p(i, j) = \sum_{k=1}^{N} \frac{1}{d_k(i, j)} \tag{4}$$

where *n* is the number of occurrences of the phrase containing the terms *i* and *j* in the same sentence and within a window of *w* terms in the profile *p*, and $d_k(i,j)$ is the distance for a given occurrence of the two terms, i.e. one plus the number of terms intervening between *i* and *j*. Thus, adjacent terms would have $d_k(i,j)=1$, and so on. While the formula above assumes a linear weight decrease over distance, we could consider other ways for the distance between terms to impact the weight computation (e.g. exponential).

As an example of the computations above, let us assume we have a single document for a given user, containing only a single sentence: "The Connection Machine models the interest, expertise and referral capabilities of each user." Then the term and phrase frequencies computed would be as in Table 2. Note that this shows only normalized frequency, not overall term weight, which would depend on the inverse document frequency factor, and is not illustrated here for simplicity.

Capturing link strength between terms allows the detection of associations between terms in a sentence, no matter their syntactic relationship; allows term association detection to cross document boundaries by following paths in the lattice; and allows precise calculation of term association strength. Instead of using the number of times

[2] This idea can be extended to phrases containing more than two words.

terms occur together, we use the totaled inverse distance between them. Thus both frequency and closeness of association are captured.

Table 2: Sample lattice built from one sentence.

	connec	machin	model	interest	expert	referr	capabl	user
connec	1	1	1/2	1/3	1/4	1/5	1/6	1/7
machin	0	1	1	1/2	1/3	1/4	1/5	1/6
model	0	0	1	1	1/2	1/3	1/4	1/5
interest	0	0	0	1	1	1/2	1/3	1/4
expert	0	0	0	0	1	1	1/2	1/3
referr	0	0	0	0	0	1	1	1/2
capabl	0	0	0	0	0	0	1	1
user	0	0	0	0	0	0	0	1

Similar ideas have been reported in information retrieval (IR) where co-occurrence statistics have been used for thesaurus construction [12] and for relevance feedback [13]. Distance-based collection and use of co-occurrence statistics have also been used at the character level for Japanese word segmentation [14]. We believe the utilization of link strength for enhancing user profiles is a new approach, integrates easily with the framework of the Connection Machine, and provides benefits over approaches that extract noun-phrases and require parsing.

5.5 Profile Updates in the Connection Machine Workflow

The initial user profiles in the system will be created based on the information provided in corporate and business databases, resumes and direct input from the users. While these initial profiles are not necessary for the system's operation, they will reduce the number of referrals needed until a suitable Provider is found during the initial phases of the application.

Once the system starts being used, the interest, expertise and referral profiles will be updated directly from the interactions as outlined in Table 3.

The user will also be able to manually update and manage his/her profiles by adding relevant documents or keywords. The profile changes caused by a user's interactions with the Connection Machine are visible in the profile section of the application as well and can be removed by users if they should choose the do so.

These updates all affect the term and link weights in one or more user's lattice-based profile. For example, the TF-IDF weights of the terms in an added document are adjusted. This could cause changes in the IDF values of terms that appear across multiple profiles. Also, if a deletion removes all information from a profile, then N, the number of profiles changes, and N also changes when a new profile is added. Adjustments such as these are not needed in typical stable document repositories.

5.6 Profile Retrieval and Ranking

As mentioned above, our approach to finding the best expert that matches a user's query is to match queries to user's profiles. We thus need to retrieve people whose expertise profiles are in some way similar to the query, and rank the profiles from most to least similar. We can also incorporate into this process the exclusion or reduction in ranking of users whose negative profiles match the query.

We base the process on the terms and phrases in the query and in the profiles. First, all profiles with terms that intersect the query terms are retrieved. This is done for computational efficiency, since profiles that do not contain any of the terms in the query are obviously irrelevant. Computation can be further optimized by sorting the returned set of profiles according to the number of terms and phrases that overlap with the query and cutting off the profiles which fall below a threshold. Then, we calculate similarity between the query and each retained profile. Recall that each term and phrase in the query carries a weight, as does each term and phrase in a profile. Based on these, we compute the similarity between a query and a profile by determining the cosine of the angle between the profile's weight vector and the query's weight vector [11]. Other similarity metrics could be used as well.

Table 3: Profile update specifications

Situation	**Interest Profile Changes (Seeker)**	**Expertise or Referral Profile Changes (Provider)**
Seeker submits a question	Add question to interest profile regardless of the outcome (i.e. whether it's answered, withdrawn, not answered).	No update
Provider refers a question	No update	No update until new provider (aka. referee) provides an answer. If referee answers question, add to referral profile the conversation up to the point of referral. If referee refers to someone else, no update.
Provider requests question clarification	No update	No update
Seeker provides question clarification	Clarification question and clarification answer are both added to interest profile.	No update
Provider supplies an answer	No update	No update until answer is accepted by Seeker
Seeker accepts the answer	No update	Add entire conversation including any clarification or (other provider's) referral comments to expertise profile.
Provider declines to answer with "Don't know the answer"	No update	Remove question from expertise profile.
Provider declines to answer with "Too busy" or "Independence conflict"	No update	No update
Provider or referrer checks "don't send similar" box after receiving a question	No update	Add question to negative expertise profile.

6 Future Work and Summary

As next steps, we are planning to utilize the user models of the Connection Machine for tasks such as targeted content distribution to interested parties, routing content to experts for verification, personalization of portals, as well as the creation of communities of interest and expertise. By analyzing interest, expertise and referral profiles for the entire organization, gap analyses could be performed and areas of concentrated expertise or interest highlighted. The continuously changing weight and link distribution of the lattice allows capturing trends in interest and expertise.

We are also interested in experimenting with different similarity metrics that take multiple profiles and feedback ratings into account and to evaluate the applicability of Case Retrieval Nets [15] for our purposes. We are also planning to look at the limited feedback mechanisms of the Connection Machine within the broader framework of a reputation system and as a means to motivate users to participate and share their knowledge [16, 17]. Other topics we consider worth pursuing are the link between social networks and expertise location [7,18] and how the information in the lattice can be interpreted with Social Network Analysis techniques [19] to determine synonyms, antonyms and value ranges.

In summary, the PricewaterhouseCoopers Connection Machine allows information seekers to enter their question in free text, finds knowledgeable colleagues, forwards the question to them, obtains the answer and sends it back to the seeker. In the course of this interaction, the Connection Machine unobtrusively updates and refines the interest, expertise and referral profiles of each user. Rather than just locating *people*, it extends the concepts of directory systems and expertise locators and acts as a virtual (adaptive) expertise provider and *answers questions*.

7 References

1. Kautz, H., Selman, B. Milewski, A., "Agent Amplified Communication", in Proceedings of the Thirteenth National Conference on Artificial Intelligence, AAAI-96, Portland, Oregon, 1996, pp. 3-9, AAAI Press, Menlo Park, 1996
2. Ackerman, M., Pipek V., Wulf V. (Eds), “Sharing Expertise: Beyond Knowledge Management”, MIT Press, Cambridge, Massachussetts and London, England, 2003
3. American Productivity & Quality Center (APQC), “Expertise Locator Systems: Finding the Answers”, Benchmarking Report, APQC Publications, Houston, Texas, 2003 (www.apqc.org)
4. Becerra-Fernandez, I., Wang T., Agha G., Sin, T., “Actor Model and Knowledge Management Systems: Social Interaction as a Framework for Knowledge Integration”, in K.D.Althoff et.al. (eds.), WM 2005 - Professional Knowledge Management, Proceedings of the Third Biennial Conference, Kaiserslautern, Germany, LNAI 3782, pp. 19-31, Springer Verlag Berlin Heidelberg 2005
5. McDonald, D. Ackerman, A. "Expertise Recommender: A Flexible Recommendation System and Architecture", in proceedings of the 2000 ACM conference on Computer Supported Cooperative Work (CSCW 2000), Philadelphia, pp. 231 - 240 , ACM Press, New York, 2000

6. Kautz, H., Selman, B. Shah, M., " ReferralWeb: Combining Social Networks and Collaborative Filtering", Communications of the ACM, vol. 40, no. 3, pp. 63-65, March 1997.
7. Thompson C., Göker M., Langley P., "A Personalized System for Conversational Recommendations", Journal of Artificial Intelligence Research (JAIR), Volume: 21 (2004), pp 393-428
8. Bridge, D., Göker, M., McGinthy, L., Smyth, B., "Case-based Recommender Systems", The Knowledge Engineering Review, 2006, Cambridge University Press, to appear.
9. Burke, R., "Hybrid Recommender Systems: Survey and Experiments". User Modeling and User-Adapted Interaction. 12(4), pages 331-370.
10. Göker, M, 'Dimensions of Personalization and their Effect on the Knowledge Containers in a CBR system', in Göker, M. (ed.) ,'Proceedings, Workshop on Case Based Reasoning and Personalization', at the 6th European Conference on Case Based Reasoning", Aberdeen, Scotland, September 4th, 2002
11. Salton, G., McGill, M.J., "Introduction to Modern Information Retrieval", McGraw-Hill, New York, 1983
12. Schütze, H. , Pedersen, J., "A cooccurrence-based thesaurus and two applications to information retrieval" Information Processing and Management, 33(3), 307-318, 1997.
13. Jing, Y., Croft, B., "An Association Thesaurus for Information Retrieval" Proceedings of {RIAO}-94, 4th International Conference ``Recherche d'Information Assistee par Ordinateur'', New York, pp. 146-160, 1994.
14. Nobesawa, S., Tsutsumi, J., Jiang, S., Sano, T., Sato, K., Nakanishi, M. , "Segmenting Sentences into Linky Strings using D-bigram statistics", COLING, Copenhagen, 1996.
15. Lenz, M., Burkhard H.D., "Case Retieval Nets: Basic ideas and extensions", in Görz G. and Hölldobler S. (Eds.), "KI-96, Advances in Artificial Intelligence", LNAI 1137, pp 227-239, Springer Verlag, Berlin, Heidelberg, 1996
16. Resnik, P., Zeckhauser, R., Friedman, E., Kuwabara, K., "Reputation Systems", Communications of the ACM, December 2000, Vol 43, No 12, pp. 45-48
17. Beenen, G., Laing, K., Wang, X., Chang, K., Frankowski, D., Resnick, P., Kraut, R., "Using Social Psychology to Motivate Contributions to Online Communities", Proceedings of the CSCW'04, November 6-10, 2004, Chicago,Ill, ACM Press, 2004
18. McDonald, D., "Recommending Collaboration with Social Networks: A Comparative Evaluation", Proceedings of CHI 2003, April 5-10, 2003, Ft. Lauderdale, ACM Press, 2003
19. Wasserman, S., Faust, K., " Social Network Analysis: Methods and Applications", Cambridge University Press, Cambridge, UK, 1994