Association for Information Systems
AIS Electronic Library (AISeL)

Wirtschaftsinformatik Proceedings 2005

Wirtschaftsinformatik

February 2005

A Tool for IT-supported Visualization and Analysis of Virtual Communication Networks in Knowledge Communities

Matthias Trier
Technical University Berlin

Follow this and additional works at: http://aisel.aisnet.org/wi2005

A Tool for IT-supported Visualization and Analysis of Virtual Communication Networks in Knowledge Communities

Matthias Trier

Technical University Berlin

Abstract: This article introduces Virtual Communities of Practice (CoP) as a valuable and complementing instrument of Knowledge Management (KM). After discussing the role and benefits of Community Software, it is identified, that although sophisticated features are available for members, the coordinating moderators still lack special software support for their complex task of monitoring and managing the expert network structure. Based on a detailed analysis of this management role, a software tool is proposed, which automatically captures the networks of expert communities within virtual discussion groups using existing communication data. Founded on theories of Network Analysis and Information Visualization, practical examples illustrate how the communication networks can be visualized and presented in a web-enabled Management Cockpit to help a Community Moderator to increase the transparency of his Community of Practice for internal members and external stakeholders.

Keywords: Communities of Practice, People Networks, Topic Networks, Analysis and Visualization, Network Theory, Software

1   Introduction

There are two basic alternatives for the implementation of Knowledge Management (KM) in an enterprise: the document- and knowledge-object oriented codification strategy and the network- and cooperation-oriented personification strategy [Hans+99]. The latter alternative primarily aims at developing networks of employees, which eventually constitute Communities of Practice (CoP), because they consist of people bound by informal relationships who share common practices [BrDu98, LaWe91]. Hence, supporting knowledge workers includes the identification of social processes and relationships that are applied to solve a business problem [Thom+01].

The increasing emphasis on the personification strategy is also observable in corporate practice, because instead of analyzing a set of related documents, employ-

ees prefer to directly contact reference persons or experts in their problem domain to gather relevant information and hence enterprises move towards pursuing this strategy [Alle00].

How Communities of Practice complement Knowledge Management instruments is demonstrated by the model shown in Figure 1. It positions the entities People, Process/Activity, Document, and Topic as the main elements of the corporate KM domain in a relationship network and highlights their relationships [Trie03b]. The primary objective of corporate KM is now to achieve transparency about the various relationships between the four entities in order to enable efficient access to the actual resources stored in the organizational knowledge structure. For example, KM Systems are employed to provide information about the interrelations between some documents but also about the relationships between instances of a process and the documents connected to it, the persons who created the documents, and the topics, which are related to the persons.

The model also illustrates the strong link between the approach of Communities of Practice and of Process Orientation. While the latter is modelling the interrelations of activities and their connection to related documents or responsible people, communities focus on the network of relations between people and their connection to topics and documents. Both disciplines are obviously partially covering the entities of the overall knowledge structure of an enterprise (Figure 1).

The importance of the community perspective for KM is also underlined in the literature, which suggests that KM Systems should much more consider the social processes between people that combine distributed knowledge into an integrated perspective instead of concentrating on classification and storage systems for knowledge 'objects' [HoLo99].
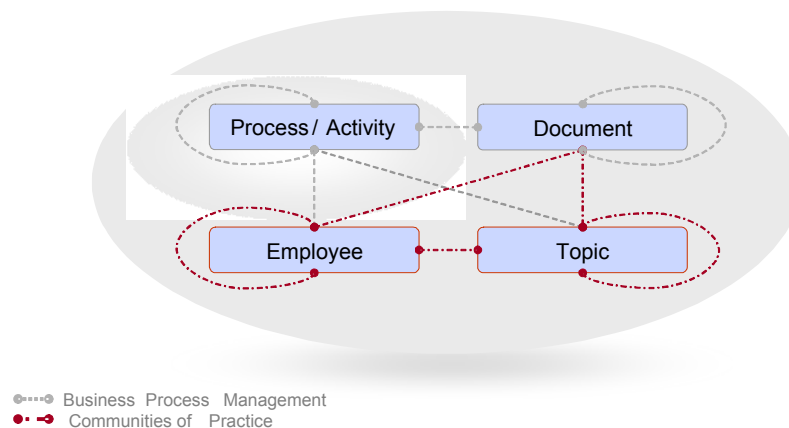

Figure 1: The Knowledge Management Entity Model [Trie03b]

2   Communities in Corporate Applications

There are various practical examples, which illustrate how enterprises approach a community-oriented Knowledge Management. These cases provide a first substantiation of the importance of moderators in corporate expert groups.

For example, Siemens is employing 'Communities of Excellence' [Enke[+]02]. There, virtual groups focus on functions, like process-engineering teams in the production or software engineers in the development division. The groups have members of the respective topical areas, processes, and projects. An IT-platform is utilized, containing discussion boards, 'urgent request' facilities, member directories, chat features, search functionality, news pages, and link collections. Official coordinators have been established and are responsible for tracking the flow of contributions to develop their subject area. Next to this individual contribution of various practical insights, members have bi-annual meetings and special community projects. In this way, the Virtual Community is enriched and backed by personal contact.

A second example for the successful application of virtual Communities of Practice is Shell International Ltd [She01]. Many informal groups already existed. They have been identified and migrated into an official global network in 1999. This network is Shell's system of Communities of Practice. Informal Groups have been officially recognized and legitimized by this procedure and hence became integrated into the enterprise's organizational structure. Shell's strategy is to have rather large groups of 1500 to 2000 people, although this must mean, that there are no real social relationships possible in such a big group. The company also established moderating responsibles, called 'hub-coordinators'. The questions and discussions mainly deal with applying a colleague's expertise to exceptional situations in the business processes, for example drilling methods. A special department is analysing all the semi-formal contributions and utilizes the insights contained to produce new process standards for the whole enterprise.

Many more corporate examples illustrate the successful application of this special instrument of Knowledge Management, e.g. ChevronTexaco Corp., BP p.l.c, IBM Corp., Unilever p.l.c, or DaimlerChrysler AG.

In such professional applications of communities in enterprises, often a coordinating role is established as an organizing and steering contact person to account for the increased responsibility of the group. Such a moderating role is also identified and characterized in the literature. Examples are McDermotts Community Leader [McDe99] or Fontaine's concept of a facilitator [Font01]. By analysing existing communities and their success factors on a more detailed level, Kim even develops a seven role model including the three roles host, event coordinator, and greeter [Kim00]. A similar concept is found by Wenger [Weng98] who even identifies seven leadership roles. Among them, there is an institutional leader who is the link to the organization, an interpersonal leader who supports tight social net-

works between group members, and a day-to-day leader, organizing activities. In the end, all these roles can also be interpreted as special organizing tasks, which can be attributed to a more general organizing role, referred to as 'Community Manager'.

Although the name of the role implies that communities are manageable, this task is very special because of the principle of voluntarism in such networks. Members dislike to be instructed and rather feel like a group of volunteers who contribute their insights to a topic only, if they need to do it. This renders management more a facilitating context management, which enables members to work on their ideas [Font01]. The generation of a strong identity and the emphasis on relationship networks is next to the content-related work a very important factor for managing such a CoP. According to this special situation, CoP managers are often emerging from the group and are equipped with strong expert legitimation to strategically and tactically be able to influence the community development.

Johnson [John01] attributes this effect to the various constructivist properties of Communities of Practice. They involve ill-structured problems, learning in a context of real-world-problems, shared goals, and the use of cognitive tools to organize knowledge. Ill-structured problems cannot be solved by any individual alone and hence the instructor is changing towards a facilitating coach for guiding the learning and helping the team develop. This moves the control away from the instructors to the group and a network of people emerges.

In addition to this special internal configuration, there are external tasks like acquiring external resources, communicating results, or connecting to other communities.

3   Supporting CoPs with IT

To design appropriate IT functionality for the coordinating roles in an expert group, the opportunities of supporting communities with software platforms have to be analysed.

CoP platforms are especially helpful for areas, in which tacit knowledge of experts can directly be applied to a related business problem [BrDu98, Weng98, WeSn00]. The people requesting help do not need to tediously analyse documents and protocols of similar scenarios to find and interpret a case with an appropriate fit to their problem. Instead, they can directly enter their request into a platform. A suitable subject matter expert can then apply his existing knowledge to this special context and does not need to explicate his experience into a broad and generic problem solution first. By answering questions of others and receiving the appropriate feedback about the practical implementation of their advice, experts are also frequently updated and reassess or even extend their experience in new concrete

application scenarios. For the initiator of the request, this method is a better way to learn by applying other's experiences.

Next to this ad-hoc mode of problem solving, Community Software provides the community with a means to discuss, develop, and integrate distributed partial approaches from projects or business processes to best practice standards. Communities of Practice are living longer than projects, which last only for a limited period. This long-term perspective of topic-oriented people networks helps the organization to maintain important competencies achieved in various related projects even after they have been completed [Weng98]. Experts generate their insights in projects and can nurture and develop their knowledge in communities. They can visually establish themselves as subject matter experts in a relevant topic field. Additionally, a valuable archive of the members' contributions is being created.

A good example for the increasing importance of IT support during the stages of the community lifecycle is British Petrol p.l.c. In the beginning, they conducted formal meetings in order to exchange expert knowledge. Next to such planned events, a large number of informal and unidentified networks existed without any rules. After the implementation of the community initiative, these groups became visible and officially recognized. The identification of these groups increased public attention and hence the relevant groups attracted more members and grew in size. Over time, the members existed in geographically very widespread locations and face-to-face contact became increasingly expensive. To compensate for the size, the communities were supported with a very sophisticated IT-platform, which provided features like mail centres, public folders, discussion boards, an integrated document storage facility, and yellow pages [McFa00].

The necessity of a central place for communication has also been substantiated theoretically by Nonaka and Teece. They established the concept of Ba, stating that knowledge transfer always requires a place like in this setting the platform in order to work. "Ba" is the Japanese word for place and represents the context in which knowledge is created, shared, systemized, and exercised [NoTe01].

In order to utilize all these advantages of software infrastructures, the manager needs to successfully migrate the very informal and invisible initial relationships of his group of experts to this platform. However, the adoption and movement to a platform has to be in line with the life cycle stage of the community [Trie02]. The expert group originally emerges from informal relationships between people, who start to develop a network without the application of information technology. Over time, the growth in group size and the geographical distribution of members directs the attention to the issue of technical support for these groups and the application of a central Community Software platform together with related service processes.

The main difficulty in employing software support is the change in network structure. A formerly decentralized network with many social elements is becoming centralized on a platform. Persons with very exclusive relationships (sometimes

established over years) could be afraid of losing their special network position. Moreover, the social character of the relationships is likely to be reduced, because IT can only support social interactions between the members of a community, but technology can rarely completely replace personal contact [Sta97] and its important contexts necessary for strong social relationships.

These adverse effects have to be compensated by the manager by means like face-to-face meetings or the establishment of a strong and visible group identity.

On a technical level, communities in an enterprise mainly develop by following one out of three migration paths [Trie03]. In the first scenario, the community platform develops from the initial application of groupware to support teams in various corporate projects. These tools are becoming modified to host defined topics and support the new user group of CoP-members. Afterwards they are offered to emerging CoPs as an internal service.

In the second scenario, the organization decides to officially align the existing expert networks and targets at connecting relevant employees without introducing a central document-centred system. When the company follows this strategy, it either develops typical CoP-functionality for internal communication and networking or it implements targeted software from a platform vendor.

In the third scenario, the enterprise already adopted the codification strategy [Hans+99] and runs a conventional primarily document-based knowledge management system (KMS). This system is being utilized by various informal groups of users. Although initially, the grouping of users is not directly intended, they form invisible communities because of their identical interests and the establishment of various relationships over time. Often companies broaden their approach towards the personification strategy [Hans+99] to directly connect their employees and reduce the problems arising from maintaining large volumes of documents, often referred to as knowledge objects. To identify and actively support the existing groups, corporate KMS's are becoming enriched by special community features for direct communication between the experts.

These multiple paths leading to IT support for expert groups already imply the heterogeneity and dynamic development of this software segment. From various related fields of applications, vendors are extending their product towards improved community support. Examples for such moving market segments are document-based knowledge bases and knowledge exchange systems, project spaces and groupware, conventional discussion boards, tools for synchronous interaction and Internet-Community software. The most widespread and conventional features of CoP platforms include discussion boards, urgent request facilities, blackboards, e-mail listservers, or membership directories. Advanced applications may additionally offer synchronous communication spaces like chats (text or video-based), document storage, evaluation systems, buddy lists, alert agents, mail centers, and calendars [Trie03].

This development towards an integrated product segment increases the risk of putting too much functionality into one platform. This can result in detrimental complexity effects affecting usage, e.g. training processes are taking longer, or researching information takes more effort. Moreover, information exchange can get inefficiently distributed over various communication channels (i.e. e-mail, discussion group, instant messaging, telephone). This segregates the expert groups into sub-groups working on the same topic but missing each other because of a preference for different communication channels.

Despite the sophisticated and very comprehensive functionality of CoP platforms, it can be recognized, that the features primarily focus community members in their communication. The coordination aspect of Communities of Practice and the according support for the responsible and moderating Community Manager is still insufficiently available. Although some simple logging functionality may be provided, it is still very unsystematic and unergonomic. Hence, the complex tasks of a moderator are not suitably supported, although this target group is in charge of important tasks like giving orientation to the group, facilitating participation, co-ordinating members and topics, or connecting the semi-structured contributions to generate insights. Appropriate software support should help the manager to answer questions like: 'How is the new topic, set up by management, accepted and developed in the group of experts?', 'Who are the central persons in the development of a specific subject field?', 'Where are the most valuable contributions and how big was the effort to produce them?', 'Who was helping others continuously without being sufficiently recognized in public?', or 'Where are inefficient parts in the expert networks that need to be actively worked on?'.

Next to the analysis of outputs, it is beneficial to provide measurements and visual insights into the actual structure of the Knowledge Community using modern and theoretically founded analysis methods. They enable a management feedback loop, consisting of observation and measurement, interpretation of measured results, and derived management interventions. This allows for a cycle of monitoring and controlling to actively manage Communities of Practice. Another factor, necessitating measurement and analysis is the requirement to report the results of community work. This is because working in an expert group is competing for time with the normal project or process work. Important stakeholders are CoP-members and sponsors (like line-, product-, top-managers).

Such a monitoring of communities is being substantiated by IBM's Watson Research Center. They concluded, that 'social translucence', which is the rich virtual impression of social structure of the communication network, is an important success factor for the effectiveness of a Community of Practice [Eric+99].

The importance of systematic monitoring and measurement has also been emphasized by the American Productivity and Quality Center APQC. Assessing the 'health' of the community has been identified as being a very important factor for knowledge management in an enterprise. Next to the incorporation of general stra-

tegic objectives of the organization and leadership qualifications of the moderating persons, the community structure is an important element of management. This institution demands, that CoPs need to set up objectives and measure the actual performance using monitoring and controlling instruments [APQC01].

Currently in corporate practice, enterprises are required to conduct survey-based audits to check the communities' status. Using questionnaires, the current conditions and outputs of the groups are estimated [Hein99]. The available data about the virtual communication is not used and integrated into this measurement approach.

Summarizing, a potential can be identified to develop software to support the monitoring and management of virtual Communities of Practice (Figure 2).
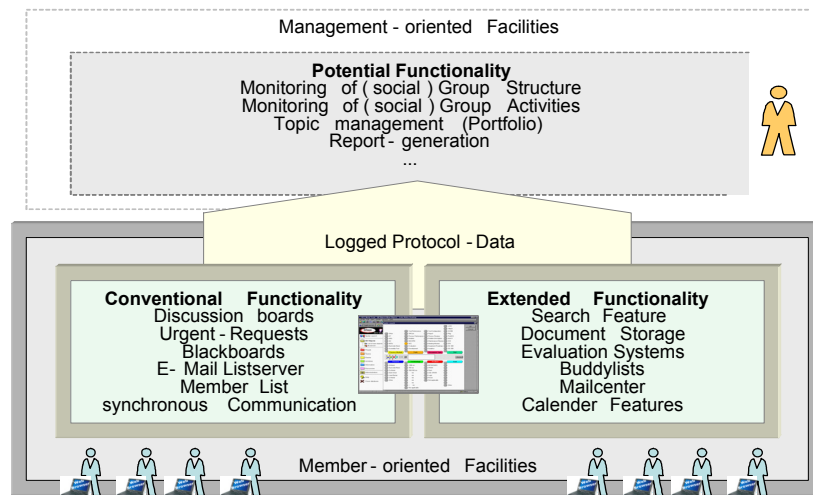

Figure 2: Adding management functionality to CoP platforms [Trie03].

4 Deriving Tool Requirements from Business Objectives

The development of appropriate software functionality for community coordinators first needs to capture the requirements of this special target group. This section now analyzes the tasks and objectives of this role. These coordinating tasks can be divided into content-related and socially motivated activities. In the next section, it is shown, how management objectives determine the definition of a useful scope for the data sets, elicited from community platforms. Subsequently, the requirements also affect the design of useful measurements and visualizations for

the communication data to finally present useful graphs of the expert group structure in a management cockpit.

In bigger Virtual Communities, direct personal contact is increasingly replaced by indirect communication via the creation and subsequent reception of written contents. Typical and additional work domains for such communities are the transformation of unstructured information, e.g. by analyzing written interactions, to quickly identify emerging information and connect this to create structured knowledge. Further activities are the diffusion of community knowledge, the use of newsletters or e-mail, the provision of relevant content from external sources, the establishment of ergonomic user interfaces, the adaptation and improvement of interaction and problem-solving structures, the generation of content for multiple reuse, the socialization of users, e.g. with membership programs, the measurement of interaction, the identification of established social relationships ('strong ties') and key persons, the connection of related persons and groups, the creation of necessary incentives, the execution of events like off-site meetings of new members with CoP-Experts, and membership management [Scho00, Par01].

Next to these content-related and transactional tasks, the literature discusses various indicators for the dominance of social motives of a Community Moderator. This includes tasks like fostering and maintaining participation with valuable feedback [LaWe91], communicating purpose, objectives, and progress [LaWe91, Hild$^+$98, NaGo98], analyzing specialization and roles of individuals to form role architectures that increase group stability [Weng$^+$02], balancing group autonomy versus openness [Pree99], creating relationship networks with tight connections and transparent visibility of members within the network [Hild$^+$98, NaGo98], establishing an environment of obligation, mutual trust, and commitment and weaken detrimental factors like concurrence and unsupportive personal profiling [NaGo98, Thom$^+$01], foster and communicate homogeneity and similarity in groups [LaMe54], or influencing orientation and objectives (polarization versus diversity) [HeWe50]. These suggestions can be extended by tasks like balancing solution exchange and solution development, creating a group identity, integrating isolated participants to improve inefficient parts of the network, monitoring the quality of interactions, sharing best practices, understanding existing informal CoP structures in order to be able to formalize a group, increasing the informal learning activities, fostering innovations, creating a familiarity between persons, or analyzing interaction and interactivity.

5   Focussing Discussion Groups

Every approach targeted at identifying patterns in network-oriented expert communication has to take into account the multiple available communication chan-

nels. They together comprise the communication network and include personal communication, phone, e-mail, instant messages, and discussion groups.

If a considerable part of interaction is based on computer-mediated communication, the expert group can be called a Virtual Community. This very widespread sub-form of a Community of Practice lends itself to further analysis, because its networks are partially visible and analyzable.

Communication via discussion groups is considered a research challenge because it is still insufficiently examined and the current interface is merely text-based. This form allows for a central and topic-oriented storage of messages between experts. Compared to this means of information exchange, the currently much researched e-mail networks have the disadvantage of being a decentralized peer-to-peer communication concept, where it is very likely to not oversee the overall content within the network. Quite contrary, discussion groups provide a consistent and complete access to the insights stored in it. The content is organized in topic threads. This makes discussion groups a suitable tool for targeted conversation generating conclusions or integrated perspectives. Examples are the development of an XML-extension to a web-based programming language, the development of an integrated design of a new business process, or the management of product problems. In all these scenarios, there are requests for expert advice in sub-domains within a larger topic area. The moderator is responsible for giving orientation and maintaining momentum within the discussing group.

On the other hand, current discussion boards are not very ergonomic. They provide features like the generation of threads. One member initiates a posting and others can reply to it. Over time a tree-like structure of comments forms around an initial question in a topic area. In larger boards, there can be thousands of semi-structured text elements posted by many hundreds of people. This makes it quite difficult to quickly work into the group's structure or to identify the most important areas and most important experts. In large groups, like the general discussions dealing with the Microsoft Operating System, the size is causing redundant contributions, so that constant analysis of the board has been implemented to identify large overlaps and cross-postings. The main reason for such inefficiencies can be seen in the user interface, which has not much changed since the first introduction of discussion board technology.

Obviously, looking at online discussions, the notion of visual components that has already been implied by the concept of social translucence as introduced in the previous section can also be applied to improve the experts' communication network: Oliver et al. find that interactive materials are essential in a virtual environment, as opposed to pure text-based scaffolding [Oliv+98]. Further, Johnson frames the question: Can Communities of Practice in their true definition be set up, maintained, and supported using current web-based applications, which are mainly text-based environments [John01]?

Following this research direction, this contribution now examines how the value creation in electronic discussions of communities can be analyzed by automatically extracting and visualizing useful and already existing data about the community structure, consisting of the entities employees, topics, and documents as well as their many relationships.

In this context, another advantage is, that the analysis of discussion groups does not cause a privacy problem like with e-mail networks, because the information contained in it is meant to be public to the members of the group. This public visibility of contributions also causes less 'noise' in the messages. This means, that in a professional application, there are almost no unrelated messages, distorting the overall conversation.

All these issues render the discussion groups a focal communication channel to further examine, analyse and visualize the exchange of knowledge in expert communities. The main objective is to make online discussions more transparent and hence easier manageable. Only then, the previously introduced requirement of regularly observing and monitoring the work of a Community of Practice becomes feasible. The following section introduces a software tool targeted at these issues.

6 Technical Concept

With the definition of management objectives and tasks and the restricted focus on virtual expert discussions, the technical architecture for a software application to support social translucence for community facilitators can be defined (Figure 3). First, the available data sources of Virtual Communities are identified, selected, and automatically accessed to retrieve the data of the communication network. These functions are provided by the application's Data Preparation Component. The data sets can then be automatically processed by the Analysis Component, using sophisticated network-oriented methods as provided by statistical and sociological approaches. They analyse author properties, thread properties, topic structures and network properties. Useful visualizations like most active or prominent authors, the acceptance level and spread of new topics, or the identification of isolated parts of the network help to generate intuitive network overviews (see section 8). These visualizations also help with the analysis of network roles or the concentration of expertise. They are then incorporated and presented in the final web-enabled Management Cockpit. Selecting, manipulating, and interpreting the visualizations and the related measurements guides management activities and helps to report important developments to stakeholders and to the members of the group.

The next sections introduce the three components together with their most innovative functionality in more detail.
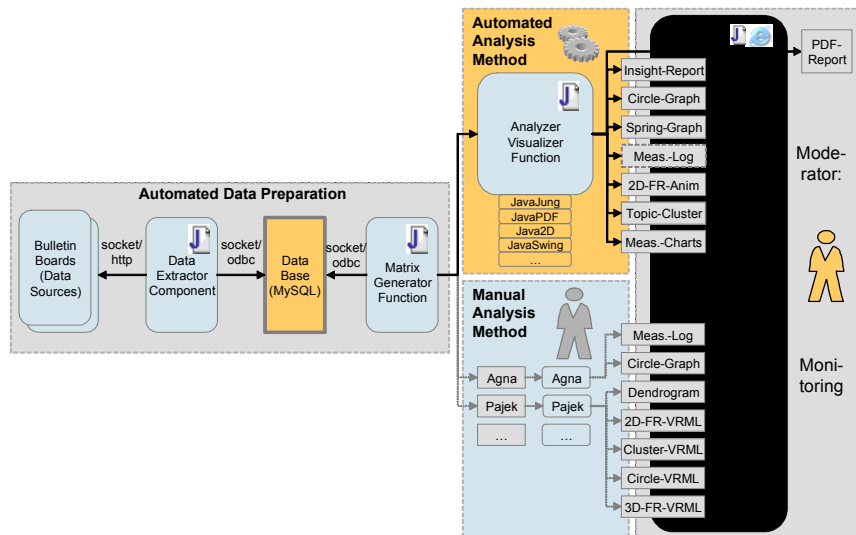
Figure 3: Technical framework and components

7 Data Elicitation and Preparation Component

The first important element of the software tool for community monitoring and management is the Data Preparation Component, which mainly includes the various Data Extraction Connectors necessary to access the virtual communication networks of selected community platforms, like Usenet Newsgroups, PHP Bulletin Board Software, and Lotus Notes Discussion Databases. These multiple connectors are necessary, because as of today, there is no standard for storing the communication network in discussion boards. However, a very widespread format to store virtual discussions is the Network News Transport Protocol (NNTP) standardized in RFC977 [KaLa86]. It is the foundation for Internet-based Newsgroups. This format defines only very few elements for storing an expert communication network on the Newsgroup Server, like a unique message identification string, the user name, the posting topic, and the posting body. Useful, but not captured are passive readers of a posting or topic keywords. To access the Newsgroup Servers, the tool's Newsgroup Data Connector establishes a socket connection on port 119. Using the RFC 977 standard, the connector sends data requests and the newsgroup server replies a variety of standardized headers, which can be parsed and interpreted to capture the necessary data.

In order to keep the succeeding tool elements independent of the diverse formats of the source platforms, the original data is transformed into a standardised data

set and stored in a MySQL-Database. The Data Connector further analyses the content of the communication using a keyword extraction algorithm. The most important keywords are then stored in the database. Later, this enables the Analysis Component to generate content-oriented analyses of the communication network.

To elicit the actual people network from the various postings, the Data Preparation Component analyses the references between the postings. They actually indicate answers or comments to a previous posting and hence a communication relation between two persons. These relations between authors are the fundamental information for creating an expert network from the data set. An example for such a hidden communication relation between two authors of a NNTP-based discussion group is shown in Figure 4.

Post-Nr.	549
Username	Anywhere
User-Adress	toff@estang.com
Post-Date	Thu, 10 Jul 2003 10:57:10 +0800
Post Topic	Apache2.0 can't start?
Post Content	...I cant'st start Apache2.0 service...
Post ID	<3f0cd63e@newsgroups.borland.com>

Post-Nr.	550
Username	Hans Karlson
User-Adress	hk@hkarl.com
Post-Date	Thu, 10 Jul 2003 13:40:43 +020
Post Topic	Re: Apache2.0 can't start?
Post Content	...put the filename between quotes...
Post ID	<3f0d50bb$1@newsgroups.borland.com>
Refers to Post ID	<3f0cd63e@newsgroups.borland.com>

Figure 4: Relation between online discussants in NNTP

After the communication data has been stored in the database, the Data Preparation Component executes the Matrix Generator Function. This element is responsible for transforming the tabular data structure stored in the database into a special matrix called sociomatrix. This is the main instrument for the sociometrical approach of Social Network Analysis (SNA), which provides a large theoretic body and collection of network measurements. One example is network density, which is the relation between the links actually present in the matrix versus the theoretically possible relationships. It shows how much of the theoretically possible communication relations are actually present. Further measures are centrality of authors, prominence, longest paths, closed loops, and various activity proxies (like in- or out-degree) for authors derived from network data [WaFa94].

8   Visualization Forms for Knowledge Networks

Based on the concept of sociomatrices for network analysis, analytical approaches concentrating on network graphs have been developed. They enable a visual analysis of large people networks. The visualization approaches actually originate in the works of Moreno in 1932 [More32], who introduced points representing actors and edges indicating the link between actors. This idea started several stages of development, like the introduction of computational procedures in the 1950's, first screen-oriented graphs of large networks in the 1970's or the event of statistical social network analysis tools in the 1990's [Free00].

The creation of such graphs from the sociomatrix generated in the Data Preparation Component is the objective of the succeeding Analysis Component. Its browser-based Java-Applets manipulate and process the matrix to generate textual analysis fragments and to render different visualizations. These results are loaded into content containers of the succeeding web-enabled Cockpit Component.

A very important such visualization element is the clustered 2D network graph using the Spring Embedder Algorithm [FrRe91]. It provides a detailed insight into the structure of a virtual discussion group. The underlying mechanism simulates a force system of virtual springs, attached between authors. In the beginning, a matrix is computed, containing the optimal distance between any two members. This distance is derived from the strength of their connection. Authors who have a strong relationship are bound by a higher attractive force and hence should have a smaller distance than authors with a weak relationship.

Then, nodes representing community members are randomly allocated on a two-dimensional plane. This results in a random actual distance. Afterwards, the complex system of springs is relaxed. The simulation compares the current with the optimal distances. The differences are stored in a force matrix. It is used to compute attractive forces that reduce a positive difference (i.e. where the actual difference is still higher than the optimal) or repulsive forces that increase a negative difference. Following [FrRe91], the formula for the repulsive force acting on a pair of nodes is -k2/d and the attractive force is equal to d2/k, where d is the distance between two nodes and k is a spring constant. By adding one node's forces towards or away from all other nodes, a final force vector can be calculated to move every node for a certain distance into the resulting direction. This process is repeated until the complete force system approaches an energy minimum. This implies that the sum of the differences between the actual and the optimal distances has been found with the current configuration of nodes and the spring system is in its most relaxed position. Thus, during multiple iterations, a clustered network graph is emerging, showing areas of strong relationships versus areas where there are no relationships.
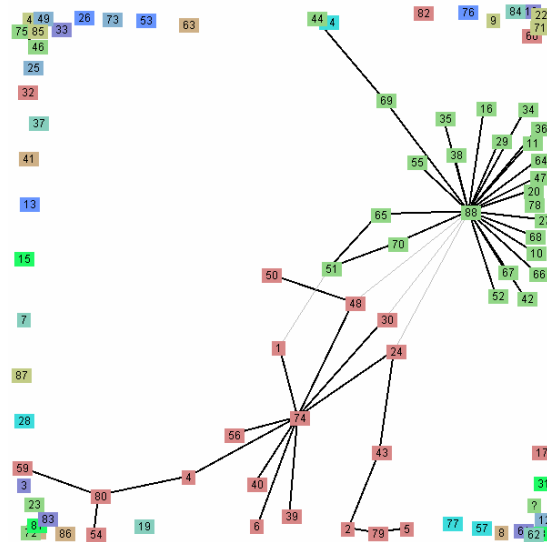
Figure 5: Two-dimensional clustered network graph of a virtual discussion

Figure 5 shows project results for a visualization of a virtual discussion group of 88 people, developing an XML-extension for the web-based programming language PHP. Isolated persons are moving to the edges of the graph. Clusters of people are visible around person 88 and 74. Obviously, they are the centre of this virtual discussion. The light relationships imply a weak relation between the two clusters. If the important bridging people were taken out of the network, it would break in two parts. For the project, this visualization has also been migrated into a three-dimensional world. This provides a more intuitive model, as it can be rotated to get a better impression of its structure. The result is shown in Figure 6. The same discussion network is visualized. To improve the transparency of this molecular-like model of this expert group, isolated people are hidden, and important authors are indicated by larger node sizes. Very intuitively, the observer can identify the two most important persons. The two experts in Figure 6 obviously do not talk to each other directly and hence build two dominating clusters around them. Still, there are four bridges between the two clusters. The four connecting people actually represent knowledge brokers, linking two sections of the network. Further it can be observed, that various people are only very indirectly linked to the network by hanging on 'tails'. They are very dependent on the person who attaches them to the network.
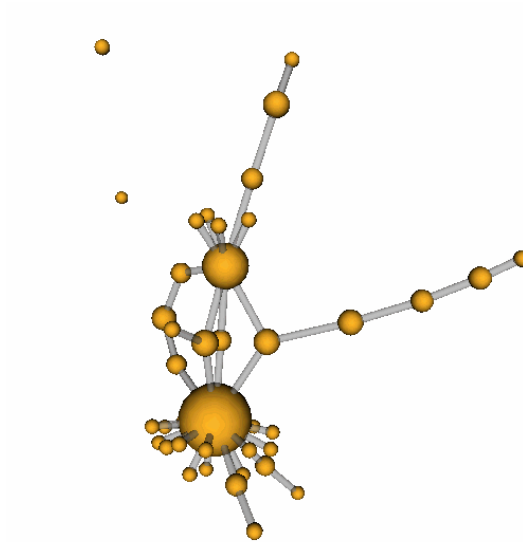
Figure 6: Three-dimensional clustered network graph of a virtual discussion

The keyword extraction algorithm of the Data Preparation Component has already been introduced in section 7. The keywords identified can also be utilized for a content-oriented analysis of the people network, like shown in the example in Figure 7. The discussion group introduced above involves the members Y, S, H, and F. The keywords extracted for the communication between Y and S have been analysed and compared. The terms 'Fopen' (indicating file manipulation) and 'HTTP' have been found in both contributions and hence are put into the middle of this relationship. The concepts 'id', 'bugs' and 'net' have only been used by Y and are hence placed nearby this node. The subnet shows that H also talked about 'Fopen' in a conversation with Y. The moderator can now search for topical concepts in his network and highlight the according subnet. For example, searching for the concept 'Fopen' results in colouring the according nodes H, S, and Y, the edges in between, and the found search terms.

This visualization gives insights about the actual topic domains and their distribution across the members of the network. The moderator can select topics and assess how they are developing. Just as in a real discussion it gets much clearer, who in the group talks about which topic. This perspective is hence much more realistic than hierarchically structured listings of texts and authors distributed across the whole discussion board.
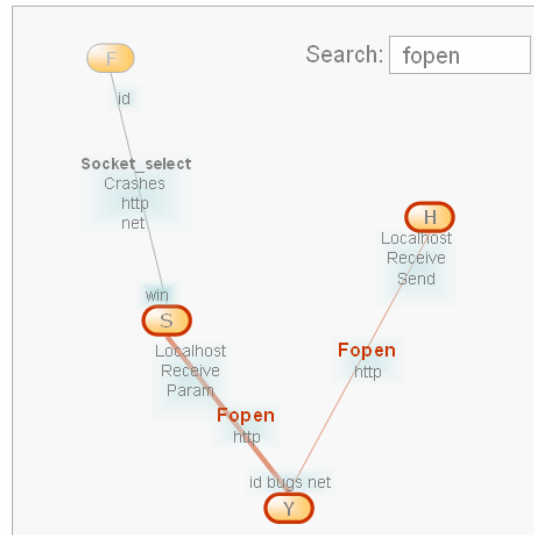
Figure 7: Integrating Topic and People Network Analysis

9 Conclusion

This contribution introduced Communities of Practice as a very well accepted approach to Knowledge Management. Its complementary position has been indicated in the literature and can be conceptualized by locating the concept of Communities of Practice in a systematic model of the relevant entities of Knowledge Management. The employment of this KM instrument in practice is implying organizational issues of coordination, monitoring, and management as well as technical issues of moving a group of experts onto a supporting CoP-platform. By comparing the tasks and objectives of the coordinating roles with the currently available IT-support for virtual Communities of Practice, it becomes obvious that this target group is not sufficiently recognized. This gap and the deficient design of current discussion group interfaces motivated the development of a module, which can be added to current platforms. It aims at eliciting the structure of the expert group by analysing the communication networks stored in data archives. Using sophisticated mathematical, physical, statistical, and sociological methods, this set of data can be visualized as an actual network of experts. Here, the integration of topic and people network analysis is regarded as the most innovative functionality. The modular visualizations are stored as complex objects and are presented in the final WebCockpit. Here, the coordinating role can consult automatically produced textual analysis and link it to the computed graphs. The manager can furthermore add

manually editable comments, like planned activities or phenomena to be documented. In the end, a detailed report can be produced, which contains selected and individually configured graphs, conveying major insights. This helps the moderator to visualize and communicate the benefits and outputs of his expert network to his group and the external stakeholders in his organization.

Future research will focus on practical applications of the Commetrix System [Comm04] to analyze existing virtual networks and further exploration and development of innovative visualization forms, including topic-people networks, evolving time-based networks, and three-dimensional graphs.

References

[AdFr00] Adams, E. C.; Freeman, C.: Community of Practice - bridging technology and knowledge assessment. Journal of Knowledge Management 1, 2000: pp. 38-44.

[Alle00] Allee, V.: Knowledge Networks and Communities of Practice. OD Practitioner Online 4, 2000.

[APQC01] APQC: Assessing the Health and Effectiveness of Communities of Practice. http://old.apqc.org/free/articles/dispArticle.cfm?ProductID=1389, 2001, Download 2003-01-20.

[Boyd+02] Boyd, D.; Lee H.-Y.; Ramage, D.; Donath, J.: Developing Legible Visualizations for Online Social Spaces. Proceedings of Hawaii International Conference on System Sciences, Big Island, Hawaii, 2002.

[BrDu97] Brown, J., Deguid, P., 1997. Organizational learning and communities-of-practice: toward a unified view of working, learning, and innovation. Organizational Science February 2, 1997, pp. 40–57.

[BrDu98] Brown J. S., Duguid P.: Organizing Knowledge. California Management Review 3, 1998: pp. 90-111.

[Comm04] URL: http://www.commetrix.de

[Enke+02] Enkel, E.; Heinold P., Hofer-Alfeis, J., Wicki, Y.: The power of communities: How to build Knowledge Management on a corporate level using a bottom-up approach. In: Davenport, T., Probst, G. (Edts.) Knowledge Management Case Book. Wiley: Erlangen, 2002, pp. 108-127.

[Eric+99] Erickson, T.; Smith, D. N.; Kellogg, W. A.; Laff, M.; Richards, J.T.: Bradner, E.: Socially Translucent Systems: Social Proxies, Persistent Conversation, and the Design of Babble. Proceedings of CHI 1999. ACM Press: New York, 1999.

[ErLa01] Erickson, T.; Laff, M. R.: The Design of the 'Babble' Time-line: A Social Proxy for Visualizing Group Activity over Time. Extended Abstracts. Proceedings of CHI 2001. ACM Press: New York, 2001.

[Font01] Fontaine, M.: Research report and toolkit: Understanding, identifying and select-ing the roles needed to staff and support communities of practice. Research Report. IBM Institute for Knowledge Management, Cambridge, 2001.

[Free00] Freeman, L.C.: Visualizing Social Networks. Journal of Social Structure 1, 2000.

[FrRe91] Fruchterman, T. M. J.; Reingold, E. M.: Graph Drawing by Force-Directed Placement. Software - Practice & Experience 11, 1991, pp. 1129-1164.

[Hans⁺99] Hansen, M.T.; Nohria, N.; Tierney, T.: What's your strategy for managing knowledge? Harvard Business Review, Mar-Apr 1999: pp. 106-116.

[Hein99] Heinold P., "The power of communities: how to build Knowledge Management on a corporate level using a bottom-up approach". Online Proceedings of Workshop MK2, Munich, 2000. http://www11.in.tum.de/forschung/foren/mkmk/proceedings/ community/heinold.ppt, Download 2003-08-20.

[HeWe50] Hemphill, J.K.; Westie, C.M.: The Measurements of Group Dimensions. Journal of Psychology 29, 1950: pp. 325–342.

[Hild⁺98] Hildreth, P.; Kimble, C.; Wright, P.: Computer Mediated Communications and Communities of Practice. Proceedings of Ethicomp'98, Erasmus University, Nieder-lande, 1998: pp. 275-286.

[HoLo99] Hoffmann, M.; Loser, K.-U.: Supporting Community Knowledge Processes with Prospective Awareness and Prospective Control. Proceedings of ECSCW'99, Workshop on Community Knowledge. Copenhagen, Denmark, 1999.

[John01] Johnson, C.M.: A survey of current research on online communities of practice. Internet and Higher Education 4, 2001, pp. 45–60.

[KaLa86] Kantor, B.; Lapsley, P.: RFC 977: Network News Transfer Protocol - A Pro-posed Standard for the Stream-Based Transmission of News. http://www.w3.org/Protocols/rfc977/ rfc977.html, 1986, Download 2004-01-12.

[Kim00] Kim, A. J.: Community Building on the Web. Peachpit Press: Berkeley, 2000.

[LaWe91] Lave, J.; Wenger, E.: Situated Learning. Legitimate Peripheral Participation. Cambridge University Press: Cambrige, 1991.

[LaMe54] Lazarsfeld, P.; Merton, R.K.: Friendship as social process: a substantive and methodological analysis. In: Berger, M., Abel, T., Page C. (Edts.) Freedom and Control in Modern Society. Octagon Books: New York, 1954, pp. 18–66.

[McDe99] McDermott, R.: Designing communities of practice: Reflecting on what we've learned. Proceedings of Communities of Practice 2001, Cambridge, 1999.

[McFa00] McLure-Wasko, M.; Faraj, S.: It is what one does: why people participate and help others in electronic communities of practice. Journal of Strategic Information Sys-tems 9, 2000: pp. 155-173.

[More32] Moreno, J. L.: Application of the Group Method to Classification. National Committee on Prisons and Prison Labor, New York, 1932.

[NaGo98] Nahapiet, J.; Ghoshal, S.: Social Capital, Intellectual Capital, and the Organizational Advantage. Academy of Management Review 2, 1998: pp. 243.

[NoTe01] Nonaka, I.; Teece, D. J.: Managing industrial knowledge - Creation, transfer and utilization. SAGE Publications Ltd.: London et al., 2001, pp. 344.

[Oliv⁺98] Oliver, R.; Omari, A.; Herrington, J.: Exploring student interactions in collaborative World Wide Web computer-based learning environments. Journal of Educational Multimedia and Hypermedia 2, 1998: pp. 263–287.

[Par01] Participate: Driving Sales Force Effectiveness Through Knowledge Sharing. White Paper, 2001.

[Pree99] Preece, J.: What happens after you get online? Information Impacts Magazine, Dec 1999.

[Scho00] Schoen, S.: Gestaltung und Unterstützung von Communities of Practice (in German). Dissertation, Technical University Munich, 2000.

[She01] Shell International Exploration and Production: Stories from the Edge: Managing Knowledge through New Ways of Working within Shell's Exploration and Production Business. White Paper, 2001.

[Sta97] Stamp, D.: Community of Practice – Learning is Social. Training is Irrelevant? In: Training, Feb 1997, pp. 35-42.

[StBr96] Steward, T. A.; Brown, V.: The invisible key to success. Fortune, Aug 1996, pp. 1-4.

[Thom⁺01] Thomas, J. C.; Kellogg, W. A.; Erickson, T.: The knowledge management puzzle: Human and social factors in knowledge management. IBM Systems Journal 4, 2001.

[Trie02] Trier, M.: Communities - Management und Unterstützung (in German). In: Gronau, N. (Edt.) Wissensmanagement - Strategien, Prozesse, Communities. Gito-Verlag: Aachen, 2002, pp.149-159.

[Trie03] Trier, M.: Anforderungen an die IT-Unterstützung für Communities of Practice (in German). Industrie Management 3, 2003: pp.49-53.

[Trie03b] Trier, M.: IT-gestütztes Management virtueller Communities of Practice (in German). Extended Abstract of DoIT-Kongress Karlsruhe 2003. http://www.integrata-stiftung.de/files/p14_trier.pdf, 2003, Download 2003-12-12.

[WaFa94] Wasserman, S.; Faust, K.: Social Network Analysis: Methods and Applications. Cambridge University Press: Cambridge, 1994.

[Weng98] Wenger, E. C.: Communities of Practice: Learning as a Social System. Systems Thinker 5, 1998.

[Weng⁺02] Wenger, E.; McDermott, R.; Snyder, W.: Cultivating Communities of Practice: A Guide to Managing Knowledge. Harvard Business School Press: Boston, 2002.

[WeSn00] Wenger, E.C.; Snyder, W.M.: Communities of Practice: The Organizational Frontier. Harvard Business Review, Jan-Feb 2000, pp. 139-145.