Community Memory Activation with Collaboration Patterns

Aldo de Moor

VUB STARLab
Pleinlaan 2
1050 Brussels, Belgium
ademoor@vub.ac.be

Abstract[1]

We present a model of collaboration patterns as reusable conceptual structures capturing essential collaboration requirements. These patterns include *goal patterns* (what is the collaboration about?), *communication patterns* (how does communication to accomplish goals take place?), *information patterns* (what content knowledge is essential to satisfy collaborative and communicative goals?), *task patterns* (what particular information patterns are needed for particular action or interaction goals?), and *meta-patterns* (what patterns are necessary to interpret, link and assess the quality of the other collaboration patterns?). We show how these patterns can be used to activate communities of practice by improving their collective, distributed memory of communicative interactions and information. We outline an approach that structures how collaboration patterns in communities of practice can be elicited, represented, analyzed, and applied. By presenting a realistic scenario, we illustrate how community memory could be activated in practice.

Introduction

Communities are living organisms that build up a collective history of information created, discussions conducted, tasks performed, and goals accomplished. Especially in virtual communities, traces of these activities can be found in the many technologies used, like mailing list archives, web pages, and document repositories. However, due to the longitudinal and distributed nature of communities, it is very hard to recollect and recombine relevant information in detail. Since memory by necessity is an abstraction of this information, conceptual patterns are potentially very useful to retrieve and link community information and trigger members of a community into active collaboration. We present a model of collaboration patterns as reusable conceptual structures capturing essential collaboration requirements. These patterns include *goal patterns* (what is the collaboration about?), *communication*

[1] The research described in this paper was partially sponsored by the EU-FP6 DIP 507483 and the EU Leonardo da Vinci CODRIVE B/04/B/F/PP-144.339 projects. This paper will be published in the Proceedings of the 3rd Prato Community Informatics Research Network Conference (CIRN 2006), Prato, Italy, October 2006.

patterns (how does communication to accomplish goals take place?), *information patterns* (what content knowledge is essential to satisfy collaborative and communicative goals?), *task patterns* (what particular information patterns are needed for particular action or interaction goals?), and *meta-patterns* (what patterns are necessary to interpret, link and assess the quality of the other collaboration patterns?). We show how formalized collaboration patterns can be used to activate communities of practice by improving their collective, distributed memory of communicative interactions and information. We outline a methodology that structures how collaboration patterns in communities of practice can be elicited, represented, analyzed, and applied. By analyzing a realistic scenario, it is shown how community memory and performance can be increased.

Collaboration Patterns

Community memory means different things to different people. It has been defined as an open-ended set of collective knowledge and shared understandings developed and maintained by the group (Marshall et al., 1995) . This focuses on the *content* of the memory. Our interest, however, is more in the *capacity* of the community members or their supporting technologies to remember relevant things about the community. This capacity includes the human minds and electronic knowledge bases storing tacit and explicit knowledge, as well as the ability to retrieve relevant information for a particular problem at hand. This memorization process should be a semi-automated process, in which tools augment people in retrieving relevant community knowledge. The strength of people is that they are very good at interpreting the relevance of a wide range of ambiguous context factors, like location, purpose, culture, stakeholder interests, history, and so on. The power of computers is that they can represent and reason about large quantities of data structures and dependencies between data items. Computer-mediated communities create vast amounts of fragmented and unstructured information over time, both as content in, for example, wikis, and as traces of conversations in, for instance, discussion fora or mailing list of archives. Much of this information is stored in human and electronic community memory. The capacity of humans by themselves to identify relevant information from this memory is very limited, however. People drown in information overload, are only vaguely aware of possibly relevant information, know about a particular item, but forgot who to contact or in which file or folder to look, and so on. Especially in collaborative communities, goal-oriented communities in which the effectiveness and efficiency of interaction processes is paramount, this collective memory loss can come at great cost and lost opportunity.

Patterns and pattern languages are getting increasingly important in community informatics (Chai & Khine, 2006; Schuler, 2002). Patterns define relatively stable solutions to recurring problems at the right level of abstraction, which means that they are concrete enough to be useful, while also sufficiently abstract to be reusable (de Moor, 2005). *Collaboration patterns* are a particular class of patterns that capture socio-technical lessons learnt in optimizing the effectiveness and efficiency of collaboration processes. As such, they can play a central role in improving community memory of collaborative communities. A typology of such patterns is given in Table 1.

Goal patterns are conceptual representations of community & individual objectives. *Examples:* A community objective can be to write a joint advisory report to be used in a policy making process. An individual objective for somebody can be to do the editing of particular sections of that report.
Communication patterns are sets of related communicative workflow and norm definitions describing acceptable and desired communicative interactions within a community. *Examples:* One possible communication pattern is that all recipients of a report must be polled for comments in an asynchronous brainstorming process. After the draft report has been released, a private, synchronous evaluation discussion must take place between the recipients, in which the list of their ideas is discussed one-by-one. If no approval can be reached in the discussion, the authors of the report must be invited into the discussion for further deliberation.
Information patterns are conceptualizations of content knowledge obtained from knowledge analysis activities. They define the knowledge elements essential for the collaborative process. *Examples:* Using advanced knowledge mining techniques, in combination with conceptual pattern analysis techniques, the reports available on an EU web site on climate change are classified in terms of types of authors, main climate issues addressed, and types of stakeholders mentioned.
Task patterns define which information patterns are to be created in particular steps in the communicative process, thus describing the role content can play in collaborative communication[2]. *Examples:* list of agenda items that (minimally) need to be produced to conclude the planning stage of a report evaluation discussion process.
Meta-patterns are conceptual patterns necessary to interpret, validate, link, and assess the quality of other collaboration patterns. *Examples:* One meta-pattern describes that information patterns used in a negotiation communication pattern may not be older than a certain date and must be subscribed to by at least two stakeholders from opposing stakeholder groups.

Table 1: A Typology of Collaboration Patterns

Collaboration patterns can be used to create a conceptual layer on top of a community information system consisting of data resources, information and communication technologies and human users. The patterns defined in this layer help *trigger* human and machine processes when certain states and events in the information system occur.

Collaboration Pattern-Based Community Memory Activation

Collaborative communities, for example knowledge management communities in corporations, e-learning communities, or research communities, are much less hierarchically structured than classical organizational forms which have well-defined roles, responsibilities, and resources. This is a strength because it allows such communities to more rapidly create new ideas and connections across networks of agents and resources. However, this lack of structure is also a weakness in that it makes the governance of the community activities much more difficult. Community governance addresses what community members may, must, or may not do, including activity design and a continuous change management process of calibrating available technologies with the social requirements (Kling et al., 2000; Preece, 2000; Shneiderman, 2002). In earlier work, we defined community governance as the regulation of community behaviour by applying community-defined norms and rules that prescribe what regulation behaviour may, must, or may not be performed by members of the community in their various roles (de Moor & Wagenvoort, 2004). In this paper, we focus on structuring one particular part of this governance: using collaboration patterns to design more sophisticated community memory activation mechanisms. We only provide an initial exploration of this novel topic. However, the examples we give outline some of the many possibilities in which collaboration pattern-driven community activation may be operationalized.

For community memory to be useful, it needs to be activated by triggering the appropriate information retrieval processes in community technologies and human members when a certain information or communication need emerges. As we have said, collaboration patterns can act as these triggers. A number of questions need to be answered to help identify the role collaboration patterns can play in community memory activation: How are these patterns represented? At what level of abstraction should they be defined? How do they come about? How can they be used?

Representing collaboration patterns

---

[2] We are indebted to our colleagues, in particular Stefanie Lindstädt, from Know-Center, Graz, Austria, for this insight.

How are collaboration patterns represented? Many different representations are possible. Initially, patterns will have a rather informal notation, such as those collected in the Pattern Language for Living Communication currently being developed[3]. However, to be able to reason about and intelligently use the numerous patterns which will be needed in a wide range of usage contexts, informal representations no longer suffice. To capture and reason about formalized patterns, *ontologies* are essential. An ontology is an explicit specification of a conceptualization (Gruber, 1994). Usually, an ontology formally represents the main concepts and relations agreed upon by key stakeholders in a particular domain. It often also includes a set of rules to reason about these representations. A very simple example of a formal representation of a goal pattern is given in Fig.1. The pattern is in conceptual graph notation (Sowa, 1984), but many other ontological representations exist.

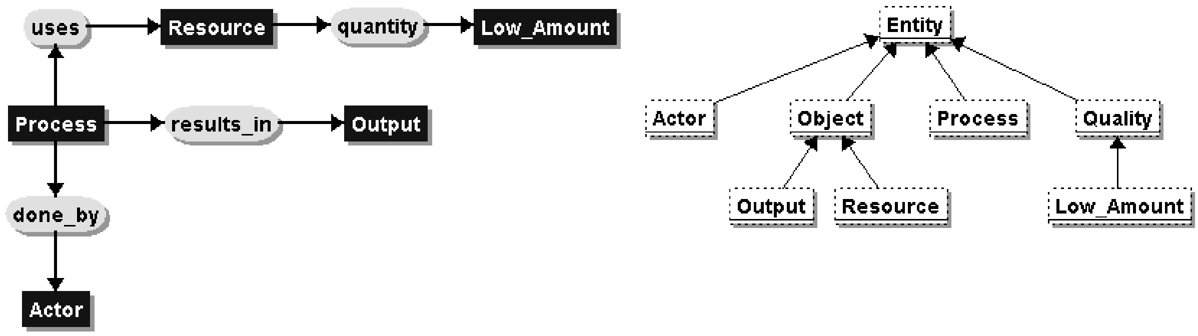

Fig. 1 A formal representation of a goal pattern

The goal pattern shows that a certain process is to be done by an actor (any role in the community), and should result in an output of some kind. To do so, one or more resources can be used, however, resources should be used sparingly, so that they do not exceed a low amount (according to whatever standard the particular community uses). On the right, the associated *concept type hierarchy* is shown. This is basically a simple meaning tree, showing that there are four types of entities: actors, objects, processes, and qualities. `Output` and `Resource` are a kind of object, `Low_Amount` a kind of quality. Many refinements of such basic patterns are possible. For example, STARLab's ontology engineering framework DOGMA allows for the addition of many different types of formal *semantic constraints*, such as frequency constraints (Meersman, 2002; Spyns, 2005). Two examples of frequency constraints could be that a minimum of one actor should be involved in the process, and a maximum of five resources should be used.

[3] http://www.cpsr.org/program/sphere/patterns/

The granularity of collaboration patterns

We have already seen that patterns should be defined at an appropriate level of abstraction or granularity, not too generic and not too specific in order to be (re)usable. One important reason is that those *accessing* community memory by necessity cannot specify their request at the level of detail of the information sought. However, also from the side of the *creators* of community memory, collaboration patterns need to be limited in amount and detail, since most people find the definition, refinement, and use of sophisticated domain descriptions to be difficult and insufficiently rewarding (Marshall et al., 1995). Thus, collaboration patterns ideally need to be specified using a number of key (ontological) domain concepts, while the role they play is to be outlined in a sound and complete conceptual model of community memory creation, representation, and use. Given that these processes are still ill-understood, such a conceptual model can only be premature. To get some general idea, in this paper we try to identify some concepts that are likely to be of interest for effective and efficient collaboration pattern specification and use by walking through a plausible scenario.

The origin of collaboration patterns

How do patterns come about? Some patterns can be generated automatically, using meta-patterns. For example, a meta-pattern could say that whenever more than **x** people use a particular tag in a folksonomy, this tag is elevated to the status of `Important-Concept`, with an associated rule that an entry on the glossary page of a community wiki is automatically created. Other patterns may come from best-practice pattern bases and are automatically selected when a community of a particular type has been identified. For example, a best-practice pattern base for knowledge management communities might include a pattern that says that a moderator-account needs to be automatically generated upon the initiation of a community, and an e-mail with a standard text subsequently sent to all initial participants that invites them to click the button '*Yes, I want to be a moderator*' on the community portal. Upon doing that, all access rights to the portal are immediately set correctly by the system. Still, many collaboration patterns will have to be defined by community members themselves, from scratch or as a refinement of generic patterns from a pattern base. Given that most community members are no knowledge engineers, great care must be taken to ensure that the interfaces for pattern definition closely resemble, are preferably even the same as, the interfaces the community uses for its daily operations. A natural moment for community members to go into this reflective effort is when the community experiences a *breakdown*, in other words when the habitual, standard, comfortable “being in the world” is interrupted (Winograd & Flores, 1986). Breakdowns occur, for example because a community has changed its way of working, but still uses now outdated tools. Although breakdowns are often a source of discomfort and conflict, they also provide opportunities for reflection and learning, often resulting in the community discovering better ways to operate and grow.

Using collaboration patterns

Collaboration patterns can be used for many different purposes: workflow modeling, service composition and integration, community evolution support, and so on. In this paper, we focus on their use for community memory activation. One way to conceptualize this use is depicted in Fig. 2.

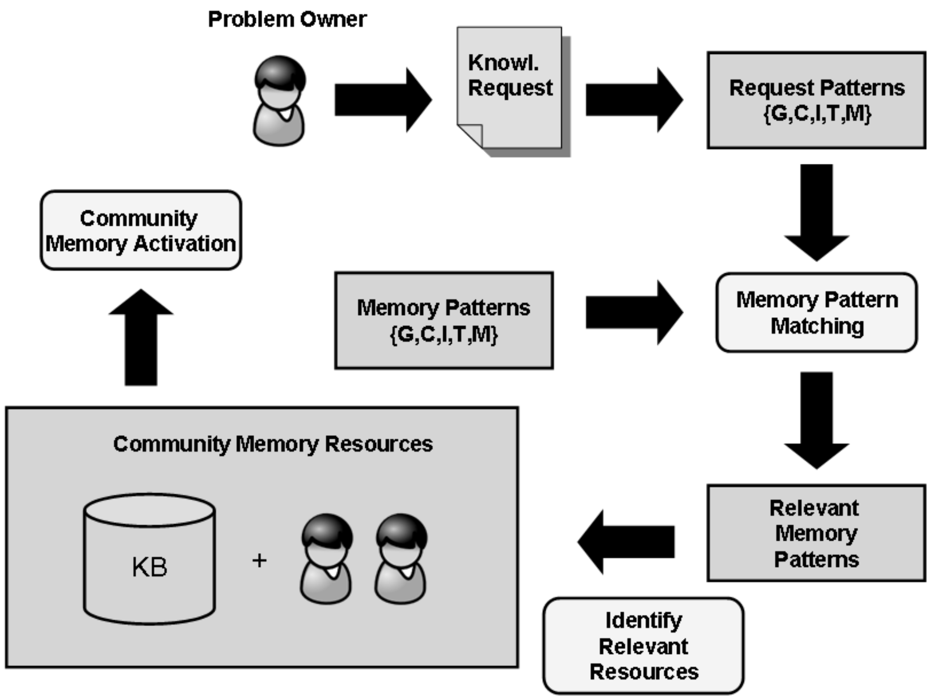

Fig. 2 Community Memory Activation Using Collaboration Patterns

A collaborative community has two main types of *community memory resources*: (computerized) *knowledge bases*[4] and *community members*. These resources are described by *memory patterns*, capturing their essential properties from a collaboration-point of view. A particular (group of) community member(s) at some point in time has a need for some knowledge which the community may provide. Their *knowledge request* is translated into a (set of) *request patterns*. In a pattern matching process the *relevant memory patterns* are selected. These are then used to identify *relevant community community memory resources*. Finally, these selected resources need to trigger some form of *(inter)action*. For example, if through the pattern matching process some moderator has been identified as the most likely candidate for answering a particular knowledge request, automatically a discussion thread could be set up in a discussion forum in which the request is posted as the first post on behalf of the problem owner, after which an e-mail is sent to the selected moderator with the link to this post and the request to reply.

A Scenario

The approach outlined above could be implemented in many different ways, to a large extent using existing technologies. It is not within the scope of this research to actually implement a system, since its main aim was to conceptualize new ways of looking at community memory, Still, to inspire community technology developers, we present a plausible scenario here, giving some hints on how these patterns could inform systems design. Collaboration patterns are again expressed in conceptual graphs notation[5].

Jane is a knowledge manager responsible for quickly setting up a task force to deal with an oil spill which has just occurred in a remote coastal area. This overall objective is captured in the following *goal pattern*:

---

[4] Of course, many other types of non-human knowledge resources exist, such as paper documents and artefacts like paintings. However, for simplicity we focus on the electronic resources prevailing in (virtual) communities.

[5] The relation type labels used in the examples are often used in conceptual graphs, and stand for: `agnt` = agent, `chrc` = characteristic, `instr` = instrument, `part` = part, `qty` = quantity, `rslt` = result.

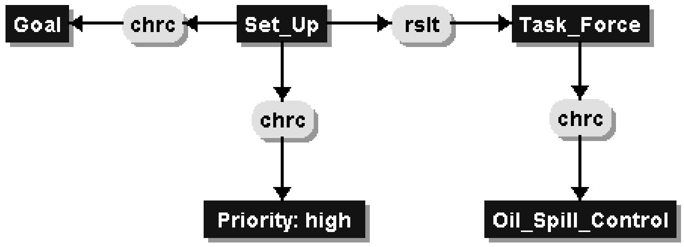

Her main *knowledge request* is to find experts who know something about a very specialized spill control technique and are willing to provide advice in a project team. Jane has little experience in this domain, but has access to a database of thousands of waste management experts worldwide, plus a long list of possibly relevant web-based discussion fora and mailing lists. It is impossible for her to examine all of these links personally, so she decides to send automatic invitations. In order to prevent spam, she also only wants to target those experts and discussion fora most likely to be interested in her appeal.

Selecting the experts from the database is relatively simple. The database contains a topic-field, which experts can fill with topics selected from a taxonomy. There are many types of oil spill control-techniques. Jane uses this simple *information pattern* to select experts from the database, which will be translated into one or more database queries selecting oil-spill topics of any type, using the notation of the database formalism used, such as SQL:

However, the pool of oil-spill experts in the database turns out to be too small to build a viable team. Jane therefore also wants to do a scan of the electronic fora most likely to contain experts to send an automatic invitation to. However, how to select the right fora from the very long list? She decides to create another *information pattern*:

This pattern could be interpreted as saying that for each discussion forum, the sum of those posts containing the string “oil spill” is to be calculated. If the quantity of that sum is larger than 50, there is a reasonable chance that the members of this forum share an expertise on oil spills. It is a crude measure, but still much preferable to selecting fora on their names alone. Of course, the pattern by itself does not do anything useful, but it can be used to configure a software agent or an intelligent discussion forum search engine.

Having selected the relevant fora with her software agent, Jane does not want to send e-mails to the complete membership of those fora. Many members will not be interested, and could consider such external requests for participation as spam. Instead, Jane only wants to send mails to those experts of the selected fora most likely to be interested. To select personalized invitations to the most relevant experts she defines this (requested) *communication pattern*:

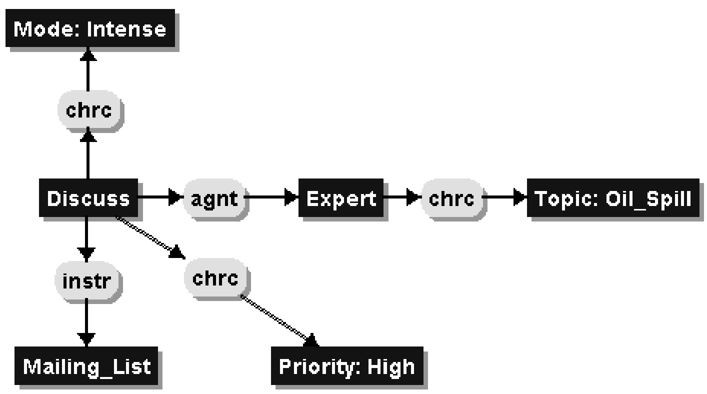

Basically, the pattern represents that the requester is looking for experts on oil spills who are willing to discuss on a mailing list in an intense mode and that the priority of the discussion is high. This pattern can now be matched with (memory) patterns representing the expertise and interests of individual experts. For example, the interests of John, an oil spill expert, might be represented by the following *communication pattern*:

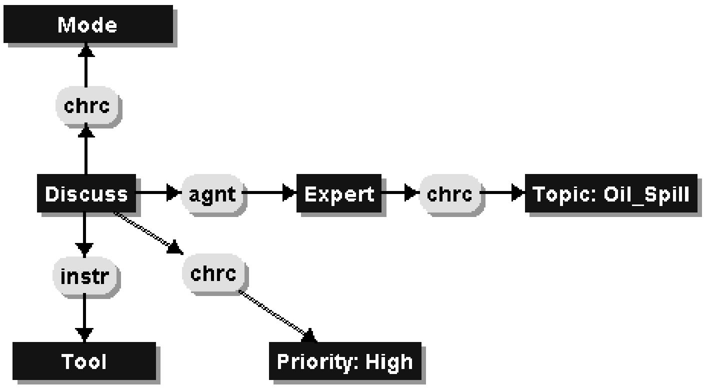

John's memory pattern is almost the same as Jane's request pattern. The difference, however, is that he does not care with what kind of tool or in what kind of mode the discussion is conducted. This pattern would therefore match with Jane's pattern[6], and John would automatically be sent an invitation to join the team. Through this sophisticated use of collaboration patterns, relevant community memory (in this case, human expertise) can be located with much more accuracy and speed than with current, brute-force information retrieval and communication approaches.

Discussion and Conclusions

In this paper, we sketched how collaboration patterns can help improve community memory by allowing for more effective and efficient use of information technologies to access human expertise and machine knowledge bases. It was only a summary sketch, and many issues still need to be worked out before collaboration pattern-driven community memory activation can become large-scale reality. Some of the issues include:

- What are the sources of patterns? The pattern knowledge is not to be provided by knowledge engineers, but by researchers and practitioners with a keen sense of the needs and capabilities of knowledge communities. Grounded theory approaches (Fernández & Lehmann, 2005; Urquhart, 2001) can be very helpful in systematically soliciting rich and legitimate patterns.

---

[6] The formal details of this matching process are less relevant here. The interested reader is referred to (de Moor, 2005) for more explanation.

- Informal patterns are more understandable to practitioners in the field, providing the nuances only verbal descriptions can convey and human beings can understand. Important initiatives are underway, like the Liberating Voices! project which aims to create pattern languages "which can help people *think about, design, develop, manage* and *use* information and communication systems that more fully meet human needs *now* — and *in the future."*[7]. However, formal pattern representation and analysis approaches are required when scaling to hundreds, thousands, or even more patterns and doing the complex reasoning needed to provide advanced semi-automated services like community memory activation. Both classes of patterns need each other. How to efficiently link them is still an open question.
- Pattern elicitation from end users is not an easy task. How to translate informal knowledge requests and offers from community members into the right request and memory patterns? Request pattern definition processes should be embedded in appropriate interfaces, for example presenting pre-set (combinations of) patterns as verbalized questions with the customization options shown as a pull-down list. Work done at STARLab on the DOGMA-MESS methodology for ontology-guided meaning evolution support focuses on exactly this problem: how to efficiently elicit and negotiate about relevant patterns and associated formal semantic constraints from evolving communities of practice? (de Moor et al., 2006)
- Pattern representation and matching processes can be operationalized in many different ways. The examples given in this paper are very simple. In real-world applications, much more complex scenarios will occur, which for example require to take into account a wide range of semantic constraints. To manage this representation and matching complexity, ontologies can play a key role. Ontologies are agreed upon representations of key concepts and relations in a domain. Much work on ontologies is being done in Semantic Web research, which has already delivered many ontological engineering approaches and tools for semantic knowledge representation and analysis. However, important issues of scalability, methodology, context, and evolution are not sufficiently addressed in mainstream ontology research (de Moor et al., 2006; Meersman, 2002). This research is now reaching a new stage with the Pragmatic Web, which focuses on how semantic technologies can actually be *used* in communities (Schoop et al., 2006). Collaboration patterns are an essential notion in developing semantic applications that are more tailored to community needs.

---

[7] http://trout.cpsr.org/program/sphere/patterns/

- Collaboration patterns are conceptual representations, but need to be activated to become really useful. Only when integrated in concrete technologies can they help realize truly active knowledge systems (Delugach, 2003), such as required for community memory activation. In this paper, we gave only some ad hoc examples of how patterns can be incorporated into concrete tools. A more systematic approach could be to think in terms of semantic (pattern-based) approaches to service-oriented architectures (Vetere & Lenzerini, 2005). Such architectures oriented more towards the social, community, than the technical, implementation-level, are still very undeveloped.

This paper applied collaboration patterns to one class of applications: community memory activation. However, a conceptual approach to reusable, best-practice collaboration pattern representation and analysis can have many more applications especially in community informatics, with (virtual) communities being so hard to describe, analyze, and support. The pattern-based approach presented here should help community (inter)action lessons to be learnt faster, and have a great impact on real communities addressing real-world problems.

References

Chai, C. S., & Khine, M. S. (2006). An Analysis of Interaction and Participation Patterns in Online Community. *Education Technology & Society*, 9(1), 250-261.

de Moor, A. (2005). Towards Ontology-Guided Design of Learning Information Systems. *Proc. of the OTM International Workshop on Ontologies, Semantics & E-Learning (WOSE'05), Agia Napa, Cyprus, November 2005.*

de Moor, A., De Leenheer, P., & Meersman, R. A. (2006). DOGMA-MESS: A Meaning Evolution Support System for Interorganizational Ontology Engineering. *Proc. of the 14th International Conference on Conceptual Structures (ICCS 2006), Aalborg, Denmark, July 17-21, 2006.*

de Moor, A., & Wagenvoort, J. (2004). Conflict Management in an Online Gaming Community. *Proc. of the Community Informatics Research Network Conference, Prato, Italy, September 29 - October 1, 2004*.

Delugach, H. (2003). Towards Building Active Knowledge Systems With Conceptual Graphs. *Proc. of the 11th Intl. Conf. on Conceptual Structures (ICCS 2003), Dresden, Germany Springer-Verlag, Heidelberg*.

Fernández, W. D., & Lehmann, H. (2005). Achieving Rigour and Relevance in Information Systems Studies: Using Grounded Theory to Investigate Organizational Cases. *The Grounded Theory Review,* 5(1), 79-107.

Gruber, T. (1994). Towards Principles for the Design of Ontologies Used for Knowledge Sharing. In N. Guarino & R. Poli (Eds.), *Formal Ontology in Conceptual Analysis and Knowledge Representation*: Kluwer.

Kling, R., McKim, G., Fortuna, J., & King, A. (2000). Scientific Collaboratories as Socio-Technical Interaction Networks: A Theoretical Approach. *2000 Americas Conference on Information Systems, August 10-13, Long Beach, CA.*

Marshall, C. C., Shipman, F. M., & McCall, R. J. (1995). Making Large-Scale Information Resources Serve Communities of Practice. *Journal of Management Information Systems,* 11(4), 65-86.

Meersman, R. A. (2002). Semantic Web and Ontologies: Playtime or Business at the Last Frontier in Computing? *Proc. of the NSF-EU Workshop on Database and Information Systems Research for Semantic Web and Enterprises, Amicalolo Falls and State Park, Georgia, USA, April 2002.*

Preece, J. (2000). *Online Communities : Designing Usability, Supporting Sociability*. Chichester ; New York: John Wiley.

Schoop, M., de Moor, A., & Dietz, J. (2006). The Pragmatic Web: A Manifesto. *Communications of the ACM,* 49(5), 75-76.

Schuler, D. (2002). A Pattern Language for Living Communication. *Participatory Design Conference (PDC'02), Malmo, Sweden, June*.

Shneiderman, B. (2002). ACM's Computing Professionals Face New Challenges. *Communications of the ACM*, 45(2), 31-34.

Sowa, J. F. (1984). *Conceptual Structures : Information Processing in Mind and Machine*. Reading, Mass.: Addison-Wesley.

Spyns, P. (2005). Object Role Modelling for Ontology Engineering in the DOGMA Framework. *OTM 2005 Workshops, Agia Napa, Cyprus, LNCS 3762*.

Urquhart, C. (2001). An Encounter with Grounded Theory: Tackling the Practical and Philosophical Issues. In E. Trauth (Ed.), *Qualitative Research in Information Systems: Issues and Trends* (pp. 114-140): Idea Group.

Vetere, G., & Lenzerini, M. (2005). Models for Semantic Interoperability in Service-Oriented Architectures. *IBM Systems Journal,* 44(4), 887-903.

Winograd, T., & Flores, F. (1986). *Understanding Computers and Cognition : a New Foundation for Design*. Norwood, N.J.: Ablex Pub. Corp.