Communities of Practice in MKM: An Extensional Model

Andrea Kohlhase[1] and Michael Kohlhase[2]

[1] Dept. of Computer Science and Mathematics, University Bremen,
kohlhase@informatik.uni-bremen.de
[2] Computer Science, International University Bremen,
m.kohlhase@iu-bremen.de

Abstract. We explore the social context of mathematical knowledge: Even though, the community of mathematicians may look homogeneous from the outside, it is actually structured into various sub-communities that differ in preferred notations, the choice of basic assumptions, or e.g. in the choice of motivating examples. We contend that we cannot manage mathematical knowledge for human recipients if we do not take these factors into account. As a basis for a future extension of MKM systems, we analyze the social context of information in terms of Communities of Practice (CoP; a concept from learning theory) and present a concrete extensional model for CoPs in mathematics.

> People don't learn to become [... mathematicians] by memorizing formulas; rather it's the implicit practices that matter most. Indeed, knowing only the explicit, mouthing the formulas, is exactly what gives an outsider away. Insiders know more. By coming to inhabit the relevant community, they get to know not just the "standard" answers, but the real questions, sensibilities, and aesthetics, and why they matter. *John Seely Brown in [Bro05]*

1 Introduction

In mathematics the production of knowledge is as dependent on social factors as in any other scientific discipline — even though this is not always realized from within, since mathematicians as a group can more easily agree or disagree on statements than other comparable groups. They use Georg Pólya's technique of "plausible deduction" that serves to differentiate between reasonable hypotheses and less reasonable ones (for a revealing ethnographic perspective on mathematics see [Hei00, 144]). Their objects of research have typically no important referent in day-to-day life, so that "truth" or "reason" is not a question of passion but of logic. At the core of mathematical identity is the concept of a proof as a process which ascertains reason [Hei00, 210]. Therefore, at first glance mathematicians build a huge, unified community and for outsiders, they seem to have the same practices all over the world. Indeed, these practices of formalization and proving can be easily distinguished from e.g. the one of experimentation by botanists. A closer look however reveals differences inside the field as well. For instance, the research objects, proof methods, proof evaluation methods, and the respective language about it differ quite dramatically even between subgroups as large as geometers and number theorists. We can discern communities of applied and pure mathematicians, which differ in the research motivation or analyticists and algebraicists, which use different mathematical tools and reasoning styles. Even on a very fine-grained level, there are communities that share or reject distinct practices, so that they can be rather small or short-lived: E.g. any research collaboration team might develop special notations (see 4.1) for their object of study and a pool of pertinent examples that are always ready at hand to test conjectures. Other examples of small, short-lived social units include the "students of a particular course".

In this paper we want to focus on the relevance of the social context for mathematical knowledge management (MKM). In particular, we want to apply the concept of "Community of Practice" (CoP) to the field of mathematics and draw consequences for the design of MKM technologies. We need to acknowledge that the context of mathematical knowledge is not only the intrinsic logical context that we model by MKM formats up to now, but also the social context. MKM can learn from this — after all, communities of mathematicians are quite efficient "mathematical knowledge management systems" and mathematicians insist that the core of mathematics lies as much in "doing" as in knowing (see e.g. [Bar02, 221]). In short — we contend that to understand mathematical knowledge management, we will have to understand its social aspects and hence to model CoPs in our systems. Otherwise we run e.g. the risk of inscribing our own CoPs into the systems, turning off users with differing practices.

2 Mathematical Communities of Practice

In 1991, Jean Lave and Etienne Wenger introduced the concept of "Communities of Practice" as the context in which learning takes place and knowledge is produced[3]. By now it is a well-established analysis tool in various fields and has experienced several extensions like "Communities of Innovation" (e.g. [Sch05, 43]) or "Communities of Knowledge" (e.g. [DP98, 66]).

2.1 Defining Communities of Practice

In order to adapt it for the field of MKM, we will now introduce the basic idea of **Communities of Practice (CoP)**, recall its definition, and argue for its relevance in MKM by interpreting Wenger's introduction of the term in [Wen99, 45][4]:

[3] They reacted with this situated learning approach against the dominant AI scheme of human intelligence as a complex computer program.

[4] In [KK05] we chose to introduce the concept via the learning object itself, mutating from raw data to information to a knowledge object within a community of practice (based on [BD00] and [PRR97]).

> *"Being alive as human beings means that we are constantly engaged in the pursuit of enterprises of all kinds, from ensuring our physical survival to seeking the most lofty pleasures. As we define these enterprises and engage in their pursuit together, we interact with each other and with the world and we tune our relations with each other and with the world accordingly, we learn. [...] Over time, this collective learning results in practices that reflect both the pursuit of our enterprises and the attendant social relations. These practices are thus the property of a kind of community created over time by the sustained pursuit of a shared enterprise. It makes sense, therefore, to call these kinds of communities communities of practice."*

2.2 Mathematical CoPs as Social Context for MKM

Unsurprisingly, mathematicians are as human as any other scientific species and as such we define and engage in common grounds, we interact and tune the relations among us and others. By doing this we produce and acquire mathematical knowledge. The interesting point that Wenger indicates here consists in his dictum "*we learn*". Even though MKM is concerned with knowledge, i.e. the product of the learning process, it seems to be interested in learning more in the form of e-learning systems as an application that be supported by MKM techniques than as a process that leads to mathematical knowledge and has to be understood for successful MKM. **Learning** is defined e.g. in Wikipedia as "*the process of acquiring knowledge, skills, attitudes, or values, through study, experience, or teaching, that causes a change of behavior that is persistent, measurable, and specified or allows an individual to formulate a new mental construct or revise a prior mental construct (conceptual knowledge such as attitudes or values)*". In this sense "**knowledge**" is a set of learned objects in an individual. Obviously, knowledge is very subjective: it depends on the learning subject. This begs the question how knowledge can become "objective" i.e. commonly accepted and understood, which is one of the central assumptions in the MKM community. In particular, how can human beings share a knowledge context? Not to drift off into philosophy, we just mention that the phenomenological concept of "intersubjectivity", i.e. the "mundane" *social* agreement on meaning, plays a decisive role in this process (for more information we suggest [Dou03, 99-126]).

Wenger continues that collective learning results in specific practices that differ with the respective community in which the knowledge was built up. Note that the term practice does not refer to a practical engagement in opposition to a theoretical engagement: "*Even when it produces theory, practice is practice*" [Wen99, 49]. Such communities of practice exist in mathematics as well (as mentioned above) even though — as they are rather informal — they don't tend to come into focus of the MKM community. We argue that in order to manage mathematical knowledge we have to pay attention to the context of production of knowledge objects. "Captured knowledge" in data bases was not only written by an author but also produced in a community of practice. Moreover, it shall be made use of by users who could be members of different communities of practice.

The supposed "common sense" or even "truth" of statements is worked out in the (social) context of CoPs[5].

2.3 What Constitutes Practice?

Etienne Wenger states that meaning must be negotiated between the interlocutors within a CoP and identifies two main, inter-operating processes in this: **participation** (action and connection) and **reification** (objectification and evaluation). In [Wen99, 63] he states "*in their complementarity, participation and reification can make up for their respective limitations*". Participation alone is too loose and confusing to establish coherent and consistent practice — therefore we e.g. take minutes in meetings. On the other hand, reified practices quickly become too inflexible to guide practice through everyday challenges hence we need to hire judges to interpret our laws.

2.4 Mathematical Practice

In MKM, we seem to have focused on just one of those processes: reification. We manage knowledge about mathematical objects via their reifications, in the same way as we have to use language objects to communicate about certain contents. At most the agreement on form can be viewed as a form of participation e.g. as valid substance equivalences. In [KK05] we present the Mathematical Knowledge Space (MKS), which we can now interpret under a CoP perspective. Wenger explicates that any "*community of practice produces abstractions, symbols, stories, terms, and concepts that reify something of that practice in a congealed form*" [Wen99, 59]. In the following we try to uncover the congealed practices.

Let us clarify this with an example: a mathematician at work. Typically, as a main part of her working life she will work by herself at her desk. She will review and evaluate what she already knows, what was said on a recent conference, or in a published journal. She might set up hypotheses and prove, postpone or drop them. Then she will elaborate on her results by writing a paper. Inbetween she will talk to colleagues, attend colloquia and conferences. Even though she works essentially by herself, she participates in the practice of mathematicians by basing her efforts on the results and values of her experience and doing it along the established (though informal) ways of the community.

2.5 Artifacts of Mathematical Practice as Living CoP Object

Clearly, authoring and studying documents are important mathematical practices. Hence, we can consider these documents as artifacts of mathematical practice. Documents result from the reification process of a practice. But similarly and at the same time documents are part of the participation process of this

[5] Remember the very controversial discussion about Hilbert's formalism versus Brouwer's intuitionism at the beginning of the 20^{th} century (see [Hei93] or [Bar02]).

practice, especially if they have been produced collaboratively. Documents are geared for the CoP-public and are a token of engagement with the specific CoP. Moreover, newcomers use these artifacts to become e.g. a mathematician, i.e. as John Seely Brown puts it in [Bro05] to "*inhabit the relevant community, they get to know not just the 'standard' answers, but the real questions, sensibilities, and aesthetics, and why they matter*". Thus, it is noticeable that reification *as well as* participation are inscribed into the working documents of the field. This property prompts us to center a CoP model around the collection of documents (in which the respective CoP practices result). In particular, we can consider it as a **living CoP object** in which many of the CoP essentials are contained and might be mined. Note that the judgmental characteristics of a CoP described by Brown blur the boundaries of such communities to be intersecting and rather fuzzy which we will make use of later on.

Now, we need to take an even closer look at the practices of a community for the modeling process.

2.6 Dynamics in Mathematical Practices

According to Wenger [Wen99, 4,49] the internal dynamics of a CoP are determined by the (interdependent) emergent characteristics of practice "meaning", "learning", "community", and "boundary", which we will now exemplify by describing them in terms of mathematical practice.

– *Practice as social negotiation of meaning:* Even though our mathematician works essentially by herself, she actively participates in the community by accepting the reifications of knowledge of her colleagues (in research documents) and sharing her own whenever possible. By this practice mathematical language can be understood among the members of the math community.
– *Practice as learning:* She works along the established (though informal) ways of the community, i.e. she reads and writes journal articles, conference papers, or listens to and speaks about colloquia talks. She takes into account the knowledge of the past in this CoP by basing her efforts on it and continuing it.
– *Practice as community:* She uses and establishes the coherence within the CoP by her engagement in the community, by working on a joint enterprise, and using a shared repertoire.
– *Practice as boundary:* She feels herself as a member of the community of mathematicians and will identify herself as such in a professional frame. But her practices will also set her apart from other communities with other common features.[6]

[6] Note that boundaries are subjective. For example, scientists in the humanities and social sciences discern CoPs of "techies" and "people people". From *that* perspective the current paper would certainly place the authors in the "techie" spot. In the MKM community we also have a division (maybe less pronounced) into corresponding CoPs. But from *this* standpoint the same paper positions the authors in the "human factor" CoP.

At the end we want to arrive at a model that is explicit enough to be implemented in MKM systems and that provides a basis to offer new services that support more of these aspects than the systems can currently offer (see section 4).

3 Modeling CoPs for MKM

In the attempt to model CoPs for Mathematical Knowledge Management we encounter the problem of a community's dynamics as the modeling process itself zeroes in on petrification. Therefore, the first question we have to deal with is: how can we get a handle on (mathematical) CoPs without inscribing the status quo — disregarding the fluid movements in a CoP? We might argue that a snapshot of the present context in the document itself is at least a first approximation. Unfortunately, this essentially yields a reduction of the idea of "living documents" to mere static (e.g. paper) documents.

3.1 Document Collections as CoP Models

Here, our approach is based on the idea that the identity of a scientific CoP is inscribed in the collection of documents this community produces. In contrast to the static characteristics of a single document, a **collection** of documents, i.e. a developing set that is structured by the respective CoP's participating members, maintains dynamic properties as the participation part of practice is involved, even embedded. From the perspective of a document in a collection, we speak e.g. of a "life cycle" of a document. Scientific developments and changing paradigms influence not only the content and form of new documents, but also the evaluation of older documents. We can even discern dynamics within a single document from this standpoint e.g. notational conventions that hold for some time but might change at any given time. Mathematical proofs serve as another example: the state of the art is not only determined by its content (search for new theorems), but also by its form (search for new proofs). If we consider a proof as combination of guarantee and explanation (see [Zin04, 4]), the explanation part is exactly the collection point of view.

In short, document collections seem to be a good starting point for modeling CoPs in MKM. Concretely, if we can assign CoP characteristics to its collection of documents, then the collection can be viewed as a *dynamic, living CoP object*, that can change without destroying the captured properties. Another advantage is that we do not need to make all properties about a collection explicit, building on the emergent effect of the composition of the documents.

3.2 Fuzzy Document Collections

Before we can start the search for CoP-characteristics in a document collection, we need to address another potential problem: CoPs do not have clear-cut *boundaries*, but a set of documents does. We take refuge in a standard idea from knowledge engineering that generalizes "sets" to "**fuzzy** sets", where set membership is generalized to a real-valued function (with values in $[0,1]$) rather than a binary predicate. Following [Zad65] we interpret the fuzzy set membership function as an evaluation function indicating the degree of membership in a CoP-defining collection, or in other words of the value of the particular document to the respective community. It seems dubious that a one-dimensional value function will suffice to express the delineation of a CoP.

Moreover, the granularity of this multidimensional value function seems to be too coarse if we only evaluate entire documents. Sometimes only specific chapters in a book, or even a single definition carry value for a specific CoP. Therefore we will identify values on “knowledge entities” or “micro-content”, i.e. document fragments that make sense as a (possibly compound) unit of knowledge.

In our approach we make use of the fact that an individual person usually does not have a crisp delineation of which documents are relevant. We will take value judgments on documents to open up the boundary of a set. In particular, we use the concept of “value judgments” to define a CoP-determining document collection as a fuzzy set. Note that this set will be fuzzy in multiple dimensions, and we will use this multi-dimensionality to support various mathematical practices in section 4.

3.3 A Multi-Dimensional Value Judgment Scheme

A natural first approach to capture such a value judgment scheme consists in using evaluation schemes that are already used regularly for peer review at conferences or journals. These should mean something for the respective community, otherwise they wouldn't ask reviewers to give feedback on these points. We will attempt to model communities of practice for mathematical knowledge management by a set of value judgments on knowledge items in documents that a community endorses.

Concretely, we will model a community of practice as a semantically closed set of documents with judgment statements on its knowledge elements[7]. A **judgment** consists of one of the following dimensions d, a reference to a knowledge element o, and a numerical value v that expresses to which degree o has dimension d.

Relevance: Is the knowledge expressed in this knowledge element relevant to the CoP?

Soundness: Are the assertions conveyed here consistent with the assumptions made by the CoP? As a special case: are they internally consistent?

Presentation: Is the presentation (not all knowledge is expressed formally) likely to be understood by the CoP members?

Originality: Does the element contain new ideas?

Significance: Will the knowledge have an impact in the community?

[7] We will specialize this when we apply our model to a concrete knowledge representation format in section 4.

The dimension of "relevance" is arguably the most important (and generic) one. It determines the mathematical knowledge endorsed by the CoP. Note that we assume that the relevance judgment is semantically closed, i.e. if an element S semantically references an element T (e.g. if S belongs to a theory that imports T), then T must be relevant to the CoP as well[8]. Note that not all properties apply to all kinds of knowledge elements, e.g. for notation declarations only the relevance property makes sense, it is used to prioritize diverse possible notations (see section 4.1). Originality is a value judgment that takes the dynamics of the knowledge creation process into account. Research-oriented CoPs usually value original ideas higher than re-iterations. The significance judgment can be used as an interim estimate or preview of actual CoP-relevance for newly contributed material.

3.4 Capturing Value Judgments for Documents

The simplest way to determine actual values for the various dimensions of a judgment consists in authoring the necessary value judgments manually; this may be suitable for an explicitly administered CoP like the aforementioned community of the students of a given lecture. Here the teacher may chose to supply not only the course materials, but also the (intended) value judgments.[9]

Another way to obtain the necessary value judgments would be to mine existing resources, e.g. from the scientific refereeing process: We already have an established process for passing value judgments on mathematical documents there. For a CoP that is centered around a particular conference (e.g. the MKM community around the annual conference on Mathematical Knowledge Management, published in this volume), we could mine the referee reports to update the CoP representation for the newly published knowledge. We imagine that referee comments would be anonymized (possibly weighted by the referee's standing in the CoP and competence). We want to point out that the judgmental dimensions above naturally coincide with those commonly used in conferences.

3.5 An Extensional Model for Mathematical CoPs

Note that the approach of identifying mathematical CoPs by the collection of documents and value judgments about it only gives us an *extensional model*, i.e. a basis for modeling certain behaviors of the community (or its members). In particular, the model does not say anything about the internal structure of the specific CoP, how membership is established or revoked, or about motivations

[8] Mathematical documents inscribe the (universal) assumption that a statement can only be accepted as reasonable, if it and all statements it depends on have been checked (i.e no "proof by authority"). Therefore we feel it is justified to inscribe semantic closure into the definition of mathematical CoPs. For other disciplines, this condition may have to be liberalized.

[9] The students may of course form a distinct CoP with their own value judgments that may or may not coincide with the teacher's.

for membership. All of these concerns are important questions[10], but currently lie outside the focus of this enterprise. Only note that since the extensional model does not address the intensional level, it does not preclude anything, and therefore may very well form the basis of future modeling efforts.

It is legitimate to ask what benefits MKM might reap from this approach, and we will answer this question in the next section by specifying some important aspects of mathematical practice that cannot currently be supported by MKM systems, since we do not have a representation of CoPs. This also shows that in the model at hand mathematical CoPs are more than mere groups of people, but a cultural phenomenon that is determined by joint practices, which determine CoP membership as a secondary aspect. Thus the model conforms to Lave and Wenger's original theory of communities of practice [Wen99].

4 Added-Value Support for Mathematical Practices

In order to evaluate, whether the implementation of the CoP model adds value to the services MKM technology can offer, we will now consider support services for mathematical practices afforded by this model from the perspectives of "meaning", "learning", "community", and "boundary" introduced above.

As we are considering concrete practices that derive from the CoP model, we need to set it in a concrete MKM representation format. The OMDoc format [Koh06] is a good basis for this, since it already contains an infrastructure for some of the mathematical practices that we want: a structured notion of theory context and an infrastructure for notation definitions. Any other format that covers these would do just as well for our purposes.

All of the applications of the CoP model we present here are related to the presentation of mathematical knowledge to humans at different levels: from notation flexibility over intra-document and inter-document discourse optimization up to the social level . That is to be expected, since the communication practices of a CoP are essential to its existence.

4.1 Meaning: CoP-Specific Notation

One of the most immediate practices in mathematics is the creation of CoP-specific languages which are represented as mathematical formulae. Their notation is one of the most visible components of mathematical documents. OMDoc represents them as objects in the OpenMath format. For instance, the equation

$$\binom{n}{k} = \frac{n!}{k!(n-k)!} \tag{1}$$

would be represented as the following string in the OpenMath XML encoding:

[10] Especially if we want to increase the participation of authors in MKM projects, see [KK04,KK05] for a discussion.

Listing 1.1. Content Markup for (1) in OpenMath

```
<OMOBJ>
  <OMA><OMS cd="relation1" name="eq"/>
    <OMA><OMS cd="combinat1" name="binomial/>
      <OMV name="n"/><OMV name="k"/>
    </OMA>
    <OMA><OMS cd="arith1" name="divide"/>
      <OMA><OMS cd="combinat1" name="factorial/><OMV name="n"/></OMA>
      <OMA><OMS cd="arith1" name="times"/>
        <OMA><OMS cd="combinat1" name="factorial/><OMV name="k"/></OMA>
        <OMA><OMS cd="combinat1" name="factorial/>
          <OMA><OMS cd="arith1" name="minus"/>
            <OMV name="n"/><OMV name="k"/>
          </OMA>
        </OMA>
      </OMA>
    </OMA>
  </OMA>
</OMOBJ>
```

Unfortunately, the machine-oriented OpenMath syntax in (1.1) is painful for humans to read, therefore it is transformed to a more human-readable form, e.g. that in (1). Note that there are varied "standard notations" for binomial coefficients: $\binom{n}{k}$, ${}_nC^k$, C^n_k, and C^k_n, that are specific to the third notation various CoPs. The third one is used by French mathematicians, whereas the last one is the Russian one.

Listing 1.2. Two Notation Declarations for Binomial Coefficients

```
<presentation for="binomial" role="applied" fixity="infix">
  <use format="TeX" lbrack="\left(" rbrack="\right)">\atop</use>
</presentation>

<presentation for="binomial" role="applied" xml:lang="fr">
  <style format="TeX">
    <text>{\cal C}^{</text>
    <recurse select="*[2]"/><text>}_{</text><recurse select="*[3]"/>
    <text>}</text>
  </style>
</presentation>
```

In OMDoc, we can define notations by embedding one of the XML fragments in Listing 1.2 into the document that defines binomial coefficients. These declarations then inform the OMDoc presentation engine which notation to generate.

Note that with a notation declaration infrastructure (see [Nay02,MLUM05] for other proposals), we can *represent* notational diversity in the OMDoc format, but not *manage* it. The attempt at a management interface manifest in the `xml:lang` attribute on the `presentation` element is a first step that allows to adapt the notation to the primary language of the document. But what about a situation, where a French mathematician writes a paper for a Russian journal, or a German professor giving a class to French students based on a Russian textbook? In such situations, OMDoc proposes to fine-tune the notation via a `class` attribute on the OpenMath `OMS` element that selects the corresponding `presentation` element.

In the model, which identifies CoPs by the collection of documents, we can generate document presentations for arbitrary CoPs instead of having to rely on target languages. Instead of embedding the notations from Listing 1.2 into the defining documents, we could have a document containing one of them in a CoP-specific notation declaration document, which would be used to generate the relevant CoP-specific style sheets used in the production of the document presentations.[11]

Note that a similar account also holds for natural language names for mathematical concepts, which behave somewhat like notations. For instance, a "ring" can be an algebraic structure for algebraicists, a subset of $\mathbb{R}^2$ that is bounded by two concentric circles for geometers and something you wear on your finger for everybody else.

4.2 Learning: CoP-Specific Discourse Building

Content-oriented MKM formats like OMDoc separate content and presentation of mathematical knowledge allowing to generate latter from the former based on general (didactic) principles and user preferences. In the last section we have seen how an explicit representation of CoPs can help manage notation choice at a formula level. At the discourse level, we also have a distinction between content and presentation: In the current view of MKM (see for instance [Far04]), mathematical knowledge is organized into a richly structured network of theories, which define mathematical objects and concepts, prove properties about them, and store examples for them. This content is then organized into discourse-level presentations — documents that contain narrative text interleaved with the content elements, and that are tailored to a particular CoP.

Applications that automate the discourse-level presentation, like the ActiveMath system [MAF$^+$01] that generates individually geared math courses, have to make choices which parts of the material to present. Here the CoP data about the micro-content together with a user's CoP membership data can be used to make informed choices. For instance, we often have multiple examples to choose from that illustrate a given construct. These will usually come from theories that are different from the theory that contains the concept to be exemplified. Of course, the examples that come from documents that are highly relevant to the reader's CoP are especially familiar and therefore have a high didactic value. The learning effect will be especially great, if a concept can be explained with an example from another CoP the reader is a member of.

4.3 Community: CoP-Specific Reference Network

Just as a we can use the CoP information to optimize the choice of material presented to a reader for intra-document discourse optimization, we can optimize the presentation of the relations of the document to the others, i.e. for inter-document discourse optimization. Each document usually refers to several other documents. If a single document is cited a lot, its importance for the discipline

[11] Thus a CoP-defining document collection is a principled resting place for notation declarations, which previously had a somewhat problematic status in OMDoc documents.

is supposed to be very high. In particular, its content is central for the CoP - its presence, its past, and its future even though the value judgment about it might shift over time. The practice of referencing reveals this shift.

Using the CoP value judgments we can e.g. weigh bibliographic references by their CoP "relevance" value. Less relevant references might be made less visible or left out altogether. To improve this further, we need to distinguish linkages in documents. In OMDoc we distinguish between semantic links (usually given as theory inclusions or theory inheritance relations), bibliographic references, and ordinary hyper-references. All of these should react differently to the CoP value judgments. To obtain the reference network data for the CoP model, we might rely on algorithms for citation relevance e.g. used in CiteSeer [cit] or GoogleScholar [Goo05].

4.4 Boundary: CoP-Specific Value Judgments

For the management of more informal CoPs we imagine to adapt techniques from *social bookmarking*, an increasingly popular way to locate, classify, rank, and share Internet resources through the use of shared lists of user-created Internet bookmarks. The reported effect of the practice of social bookmarking consists in the engaging latitude of individual tagging and the socially informed inference drawing process based on mass data. At this point we don't want to judge the pros and contras of this approach, we just draw on the emergent dynamics of this new technology (see e.g. [Wei05] for a discussion). If such tags are visualized e.g. via tag clouds then a user obtains a feeling of belonging and participation, i.e. we can integrate the boundary effects for MKM technologies.

We envision that (in analogy to general social tagging systems like "delicio.us" for documents or "flickr" for photos, or scientific ones like "Connotea") CoP members store value judgments of knowledge items and make these lists publicly accessible. The value judgments necessary for representing a CoP can then simply be computed from the harvested value judgment lists of the members. This way the developmental cycles of a CoP are mirrored in the CoP model as well.

5 Conclusion, Related and Further Work

We have argued that the MKM community is still turning a blind eye towards the "*social life*" of knowledge and is thus missing out on valuable chances to offer personalized added-value services: mathematical knowledge does not live in a social vacuum, and neglecting that will rob MKM systems of the flexibility to scale up to larger and thus more diverse user communities. We cast the discussion in terms of "Communities of Practice", which we adapt to the context of mathematical knowledge, and propose a simple extensional model that is very well-integrated into MKM practice. We have shown the usefulness of this model by exhibiting knowledge management applications that feed on CoPs thus modeled and use the information to tailor the presentation process to the (CoP of the) user.

Unfortunately, we have not yet covered a very important practice in mathematics from the CoP angle of view: mathematical proofs. Of course, mathematical documents contain proofs, and as such they are encompassed by our model, but questions like proof style, or their peculiar status as a communicative acts between guarantee and explanation (see [Zin04, 4]) will merit further study.

Our work here is related to user modeling and community building work in other MKM systems. For instance, the OMDoc-based ActiveMath [MAF+01] system employs a user modeling component to infer the prior knowledge of a reader and employs it for the user-adaptive generation of narrative documents leading up to a chosen concept. Compared with our model, the user model is more detailed (it contains graded assumptions about "knowledge", "understanding" and "application"), and less comprehensive (for instance it does not contain information about notation preferences and "user models" for groups of users are not envisioned). But the ActiveMath user model could be viewed as the representation of a one-person CoP, e.g. by interpreting the "knowledge" property as "relevance". On the other hand, CoPs could be used to prime the user model in ActiveMath by assuming that a CoP-member knows, understands and can apply concepts proportional to their relevance. This would allow the user to simply identify her (pre-existing) CoP rather than giving confidence values for a large set of mathematical concepts.

Some MKM systems try to increase author involvement by providing community features (as e.g. this is what drives the runaway success of the WikiPedia). For instance the Connexions project [CNX06] has recently added "community pages" to their namesake system [CNX05] and plans to use a community-driven post-publication system called "lenses" for quality assurance. The former offer communities a forum for discussions and a way to identify a collection of relevant course modules. The lenses in Connexions are rating systems allow communities and institutions to endorse certain course modules. However, as the CNXML system can model less "practices" than OMDoc, its reach is limited to document selection.

The next step in our research enterprise will be to implement the CoP model presented in Section 3 and the added-value services sketched in Section 4 in an MKM system. We hope that by offering added-value services we will entice users to enter value judgments that can be used to represent and identify CoPs (extensionally). A feedback/rating system with an interface like the ones used in `amazon.com`, `Slashdot` or `ebay`, could turn an MKM system into a data collection tool for studying CoPs of mathematicians — although the necessary preselection of mathematicians who are willing to use MKM systems will probably introduce a strong bias.

To this effect we are currently extending the Connexions system to cope with OMDoc knowledge. We plan to build on its existing community features and extend them with this CoP model. As the system is used quite heavily for E-Learning in diverse communities ranging from music theory to electrical engineering, we hope to gain valuable insights into the inner workings of CoPs and their relation to value judgments. Another direction we would like to pursue is the extension of the OMDoc format that it can represent more mathematical practices, not just notation declarations, and the specification of technical terms.

All in all, we view the requirements coming from CoPs as essential guidelines for the further development of MKM formats. From an analysis concerning the relations between knowledge and practice, we can deduce the relevance of mathematical practices for MKM technologies as Osterlund and Carlile conclude that "*the relational core of a knowledge sharing theory easily falters. [...] We end up instead with a perspective that focuses on the storage and retrieval of explicit knowledge represented in information systems. Knowledge becomes an object shared within and across community boundaries* **without consequence for the community in which it originated**." [OC03, 18].

References

[Bar02] John D. Barrow. *Ein Himmel voller Zahlen. Auf den Spuren mathematischer Wahrheiten.* Rowohlt, 2002. Original publication 1992 "Pi in the Sky".

[BD00] John Seely Brown and Paul Duguid. *The Social Life of Information.* Harvard Business School Press, 2000.

[Bro05] John Seely Brown. Learning in the Digital Age. online, 2005. seen on 2006-02-07 at `http://www.johnseelybrown.com/learning\_in\_digital\_age-aspen.pdf`.

[cit] CiteSeer.ist scientific literature digital library. web page at `http://citeseer.ist.psu.edu`. viewed March 2006.

[CNX05] Connexions. Project home page at `http://cnx.rice.edu/`, seen January 2005.

[CNX06] Connexions: Sharing knowledge and building communities. White paper at `http://cnx.org/aboutus/publications/ConnexionsWhitePaper.pdf`, February 2006.

[Dou03] Paul Dourish. *Where the Action Is: The Foundations of Embodied Interaction.* MIT Press, 2003.

[DP98] Thomas H. Davenport and Laurence Prusak. *Working Knowledge.* Harvard Business School Press, 2000 edition, 1998.

[Far04] William Farmer. Mkm: A new interdisciplinary field of research. *Bulletin of the ACM Special Interest Group on Symbolic and Automated Mathematics (SIGSAM)*, 38:47–52, 2004.

[Goo05] Google scholar. Web page at `http://scholar.google.com`, viewed January2005.

[Hei93] Bettina Heintz. *Die Herrschaft der Regel: Zur Grundlagengeschichte des Computers.* Campus Verlag, 1993.

[Hei00] Bettina Heintz. *Die Innenwelt der Mathematik. Zur Kultur und Praxis einer beweisenden Disziplin.* Springer-Verlag Wien, 2000.

[KK04] Andrea Kohlhase and Michael Kohlhase. CPoint: Dissolving the author's dilemma. In Andrea Asperti, Grzegorz Bancerek, and Andrej Trybulec, editors, *Mathematical Knowledge Management, MKM'04*, number 3119 in LNCS, pages 175–189. Springer Verlag, 2004.

[KK05] Andrea Kohlhase and Michael Kohlhase. An exploration in the space of mathematical knowledge. In Kohlhase [Koh05], pages 17–32.

[Koh05] Michael Kohlhase, editor. *Mathematical Knowledge Management, MKM'05*, number 3863 in LNAI. Springer Verlag, 2005.

[Koh06] Michael Kohlhase. OMDoc *An open markup format for mathematical documents (Version 1.2)*. LNAI. Springer Verlag, 2006. to appear, manuscript at `http://www.mathweb.org/omdoc/pubs/omdoc1.2.pdf`.

[MAF+01] E. Melis, J. Buedenbender E. Andres, Adrian Frischauf, G. Goguadze, P. Libbrecht, M. Pollet, and C. Ullrich. The ACTIVEMATH learning environment. *Artificial Intelligence and Education*, 12(4), winter 2001 2001.

[MLUM05] Shahid Manzoor, Paul Libbrecht, Carsten Ullrich, and Erica Melis Ma. Authoring presentation for OPENMATH. In Kohlhase [Koh05], pages 33–48.

[Nay02] William Naylor. Mappings between presentation markup and semantic markup for variable-sized objects. In *Second International Conference on MathML and Technologies for Math on the Web*, Chicago, USA, 2002.

[OC03] Carsten Osterlund and Paul Carlile. How Practice Matters: A Relational View of Knowledge Sharing. In Marleen Huysmann, Etienne Wenger, and Volker Wulf, editors, *Communities and Technologies*. Kluwer Academic Publishers, 2003.

[PRR97] G. Probst, St. Raub, and Kai Romhardt. *Wissen managen*. Gabler Verlag, 4 (2003) edition, 1997.

[Sch05] Tim Schloen. Expertennetzwerke als Innovationsschmiede - das Konzept der Communities of Innovation. In Sylke Ernst, Jasmin Warwas, and Edit Kirsch-Auwärter, editors, *wissenstransform*, pages 40–53. LIT Verlag, 2005.

[Wei05] David Weinberger. Tagging and Why it Matters. online at `http://cyber.law.harvard.edu/home/uploads/507/07-WhyTaggingMatters.pdf` seen 2006-05-22, May 2005.

[Wen99] Etienne Wenger. *Communities of Practice: Learning, Meaning, and Identiy*. Cambrigdge University Press, 1999.

[Zad65] Lotfi A. Zadeh. Fuzzy sets. *Information and Control*, 8:338–353, 1965.

[Zin04] Claus Zinn. *UnderstandingInformalMathematicalDiscourse*, volume 37 of *Arbeitsberichte des Instituts für Informatik*. Universitä Erlangen-Nürnberg, Institut für Informatik, September 2004.