Improving mathematical knowledge items by acting on issue-based community feedback

Christoph Lange[1], Tuukka Hastrup[2], and Stéphane Corlosquet[2]

[1] Computer Science, Jacobs University Bremen
ch.lange@jacobs-university.de
[2] Digital Enterprise Research Institute, NUI Galway
{tuukka.hastrup,stephane.corlosquet}@deri.org

Abstract. In informal community-driven knowledge collections like wikis, there is no well-defined way of reporting issues with knowledge items. When something is wrong or needs improvement, there is hardly any support for the community to communicate this in a focused way. The workflow of discussing about an issue, coming up with ideas on how to solve it, agreeing on the best idea, and finally putting this idea into practice in a retraceable way is sometimes standardised in terms of best practices and social conventions, but not supported by of knowledge management. We present an approach to improving this in a semantic wiki, where not only the articles contain structured knowledge, but also the discussions *about* this knowledge are structured using an argumentation ontology. We show how, by domain-specific extensions of this ontology, the wiki system can not only support a focused discussion about issues but also assist with putting solutions approved by the community into practice in many common cases. This is demonstrated on the prototype of a system in a use case from mathematical knowledge management.

1 Introduction

This article deals with *issues* with knowledge items[3] in community-driven knowledge management – how to report them, how to argue about them, and how to resolve them. An issue with one such knowledge item could, e. g., be that it is hard to understand or consists of wrong facts, that redundancies with other knowledge items are identified, or that a subpart of one knowledge item is considered to deal with a topic of special interest, deserving to be promoted to a knowledge item of its own. Members of the community report issues and argue about them, proposing solutions that are again subject to discussion, until finally a solution is implemented. Such a discourse can be lengthy and hard to keep focused, as issues can be "wicked problems", exposing traits like not allowing for a "defininitive formulation", having solutions that are "not true-or-false but good-or-bad", and the nonexistence of an "immediate and [. . .] ultimate test of a solution" [31]. A

[3] We will henceforth refer to a piece of knowledge about any distinct subject of interest as a "knowledge item", regardless of its degree of formalisation. We use the term "issue" as in "issue tracking systems".

solution is usually materialised in an improved version of the affected knowledge item or a new knowledge item. Later, other users who want to understand why some knowledge item is modelled in a particular way (consider new members of the community!), would trace back the discourse that led to its creation or modification. Thus, the discussions about issues with knowledge items become part of the collective experience of the community.

We investigate a model that allows for structured argumentation about issues with knowledge items and present a system that can assist users with the implementation of solutions in common cases. We try to keep our model as general as possible but particularly investigate it in the setting of a wiki for mathematical knowledge. In a wiki, one page usually holds knowledge about one distinct topic, or about a set of closely related topics [9]. Some typical knowledge items in mathematics are definitions of symbols or concepts, theorems, and proofs. Issues with them can be that a knowledge item is wrong, incomprehensible, presented in an uncommon style, or redundant (cf. section 4).

We continue with a motivation why issue handling is insufficient in existing wikis, an introduction to semantic wikis and argumentation ontologies, which we consider helpful in that regard, a theoretical outline of our approach, a description of a prototypical implementation, and then walk through a typical use case to demonstrate the approach in our implementation. We conclude with a discussion of the coverage of our support for argumentation, a comparison with related work, and ideas how the presented work can support communities of practice.

2 Motivation: Wikis and Wikipedia

In collaborative knowledge bases without semantic structures, such as conventional wikis, there is no well-defined way of handling issues with knowledge items and hardly any technical support. The community is left alone with establishing a manual workflow of reporting, discussing, and solving issues and documenting the solutions. This is mostly done by jointly agreeing on best practices in conflict resolution and authoring and making them official policies for the community [14]. As a concrete example, consider a Wikipedia article that violates the Wikipedia principle of a neutral point of view [41][4]. Some author who is concerned about this can tag the article by inserting the building block "Neutralität" (neutrality), which shows an according warning message [39]. It is then recommended to justify why the neutrality of the article is debated by adding a respective section to the discussion page of the article. Within that section, the general conventions for discussions pages apply [40]: The author has to make clear what section of the article his discussion post applies to, he has to verbalise his report in a comprehensible way, and finally has to append his signature (a link to his user profile with a timestamp). An author who wants to discuss an existing issue has

[4] As the audience of this article is mainly German and the different language editions of Wikipedia have developed slightly different conventions, we refer to the German Wikipedia. Pointers to other Wikipedias can be found on the respective pages by following the links to other languages ("Andere Sprachen").

to look up the corresponding section on the disscussion page and then indent his reply by one more level than the post he is replying to. Solutions to issues would also be proposed in natural language only, and if users come to vote on proposals, they would do it in an ad hoc manner, e. g. using list items prefixed with "yes" or "no". A solution for restoring the neutrality of a controversial article could be citing reliable arguments in favour of the view that has been less represented so far. Eventually, one author who is trusted by the community[5] would judge whether there is a consensus about a particular solutions, or simply count the votes, and then implement the solution approved by the community, again without any assistance from the system. A justification for the resulting revision of a page can be given by a descriptive editing comment that links to the section of the discussion page where the respective issue was discussed [42]. However, authors do not always do this, which sometimes makes it hard to retrace decisions.

In large communities like Wikipedia, these procedures work sufficiently; indeed, the quality of Wikipedia articles has been found to strongly correlate with the number of authors [4]. We are aiming at smaller communities but support them by intelligent issue tracking: a system that has a basic understanding of what types of issues are reported with what types of knowledge items, what solution are proposed, and whether people agree or disagree with these proposals.

3 Foundations: Semantic Wiki and Argumentation

In a semantic wiki, the knowledge is more structured than in a non-semantic one. Other than employing a knowledge representation based on semantic web technologies such as RDF [36] or ontologies (see e. g. [10]), semantic wikis are diverse: In some systems, shallowly annotated text prevails, whereas in others, unstructured text only appears in comments or labels that describe formal concepts. Even others mix annotated text and highly formalised problem-solving knowledge [1]. The most common approach is, however, to represent knowledge about one subject of interest – a "knowledge item", in our terminology, – by one wiki page and to annotate pages and links between pages with types defined in an ontology [28, 6]. In this kind of semantic wikis it is advisable to keep pages small and refactor them if they tend to describe more than one knowledge item. The graph of typed nodes (= pages) and edges (= links) is commonly represented in RDF [36]. Existing ontologies, such as Friend of a Friend (FOAF [11]), are either preloaded in the wiki or imported later, or a new, custom ontology is built collaboratively during the annotation of the wiki pages.

To the best of our knowledge, only semantic wiki features semantically structured discussions so far. In IkeWiki [32], the relationship between the knowledge item represented by a page and the discussion about it is represented in the RDF graph, and the associated discussion page itself is not an opaque block of text, but a self-contained discussion forum. It consists of threads and posts,

[5] Note that in the remainder of this article we will assume well-behaved and cooperative users. Encouraging or enforcing orderly behaviour is an interesting research question in itself but not considered here.

with RDF links to the user profiles of their authors. This is achieved using the SIOC ontology (Semantically Interlinked Online Communities [5]), which models generic aspects of various forms of discussion forums. SIOC itself does not yet allow for typing discussion posts as issues, proposed solutions, agreements, etc., but is meant to be extended by dedicated ontologies that give more specific types to discussion posts [3, sect. 4.2]. Indeed, making every post a distinct RDF resource and preserving the threaded structure of a discussion serves as a basis for adding an argumentative layer, as we will show in the following.

An early approach at formalising argumentation about issues was IBIS (Issue-Based Information System [17]), which particularly aimed at wicked problems [31]. The DILIGENT argumentation ontology was conceived in the context of the ontology engineering methodology of the same name as an extension of IBIS that makes arguments more focused, thus making design decisions more traceable and allowing for inconsistent argumentations to be detected [33, 34]. We found the DILIGENT vocabulary suitable for our setting of structured knowledge engineering. But note that our understanding of an *issue* differs from the DILIGENT methodology: There, issues are issues of conceptualising new aspects of a domain, and ideas refer to the formalisation of these conceptualisations, according to the well-known definition of an ontology [12], whereas in our case issues are issues with more or less formalised knowledge items that already exist.

A discourse in terms of the DILIGENT argumentation ontology is structured as follows: When an issue has been raised, community members can express their agreement or disagreement with it, i. e. whether they consider this issue important, justified, and legitimate. An issue can be resolved by implementing a proposed and – again by posting agreements – approved idea in the space of knowledge items (called "ontology entities" in DILIGENT) and concluding the discussion thread with an explanation of the decision taken. This decision will link to the issue that has been solved and to the idea i that was realised. If that idea was to create or modify a knowledge item k, a link "i resolves into k" will be created. Besides merely agreeing or disagreeing with an issue or idea, the community can also *argue* about it, i. e. users can justify it by examples or evaluations, or challenge it by alternative proposals or counter-examples, and others can again agree or disagree with these arguments.

4 Approach: Domain-Specific Ontology and Assistance

We assume that every knowledge item has at least one principal type and that these principal types are disjoint with each other. Concrete types of knowledge items will heavily depend on the application domain; in this article, we focus on mathematics and consider mathematical statements as atomic knowledge items, a statement being a definition, theorem, proof etc. We take these types from the OMDoc ontology [18, 21] that we modeled for representing mathematical knowledge originating from documents in the semantic markup language OMDoc [16] on the semantic web. This ontology abstracts from the syntax of the OMDoc language and aims at formally modeling aspects of its semantics that cannot be expressed in an XML schema (e. g. that the target of a link from a proof to the resource that is proven must be a theorem, and that certain links between knowledge items constitute dependencies) but on the other hand are too informal for being meaningful to e. g. an automated theorem prover.

Besides its principal type, a knowledge item can have additional types that will not be taken into account for reporting and resolving issues. We do not restrict the nature of these types; they could e. g. describe the status of a knowledge item in terms of project management (such as "draft", "under review", "published") or in the topics the item pertains to (such as "algebra", "statistics", "physics").

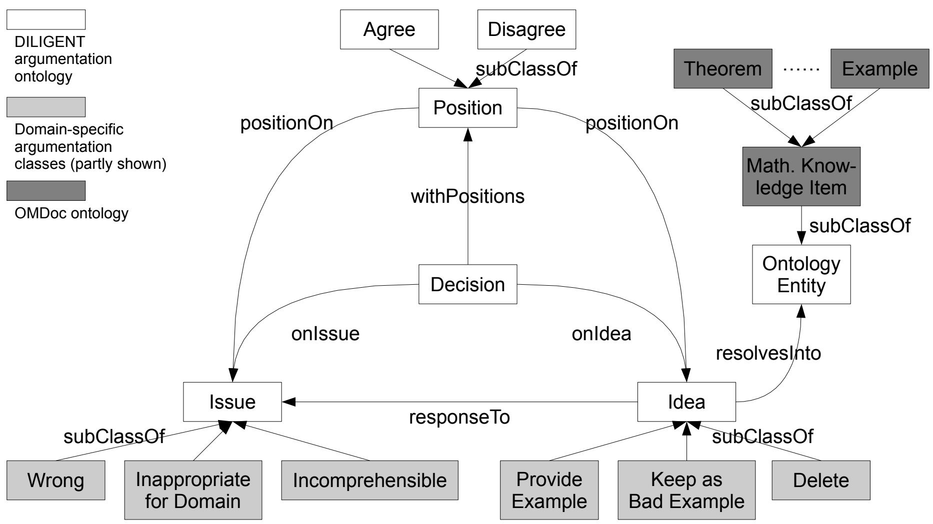

Fig. 1. The DILIGENT/OMDoc argumentation ontology

We assume that there can be a threaded, SIOC-structured discussion about every knowledge item and that the discussion posts can be typed as instances of classes from the DILIGENT argumentation ontology. In our current model, we consider the classes *Issue*, *Idea*, *Agree*, *Disagree*, and *Decision* (fig. 1). To establish a bridge between the domain of knowledge and the argumentation about it, we created an ontology of domain-specific subclasses of DILIGENT's *Issue* and *Idea* class (fig. 1). A particular type of issue is considered applicable to certain types of knowledge items; we model this in the ontology as well. For example, an issue with a mathematical proof can be that it is wrong, whereas the notation of a symbol cannot be wrong but inappropriate, misleading, or hard to write. Furthermore, we assume that to a pair of a knowledge item type and an issue type, certain types of ideas can be applied. For example, if a proof is wrong, it could be deleted and replaced by a correct proof, or it could be kept as an instructive bad example. Obviously, we do not expect to cover *all* possible cases with a finite set of predefined issue and idea types, but the most common ones.

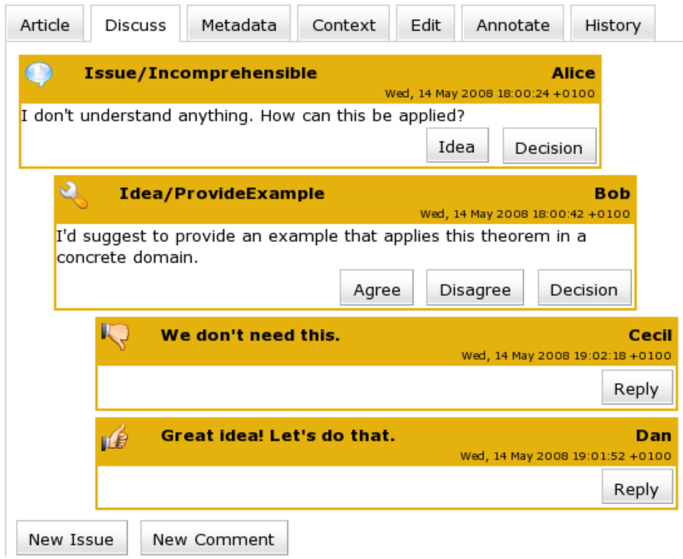

Fig. 2. In the middle of a discourse

In order to get an understanding of common issues and solution practices in the domain of mathematical knowledge management, we are currently conducting a survey among domain experts[6]. We collect information about the previous experience of the participants with mathematical knowledge bases, the support for tracking and solving issues in the tools they have used, types of knowledge items they have dealt with, types of issues they have encountered, how these issues were solved, and reasons why issues remained unsolved. So far, 42 people have participated, 20 of which have answered all questions. A majority is experienced in contributing to libraries of software tools like automated theorem provers, but many participants have also contributed to wiki-like knowledge bases. The most common granularity of knowledge items they have experienced is either a course unit, a mathematical theory (i. e. a few related definitions and axioms), or a mathematical statement. Only in a few cases there was support for automated issue tracking and solving. The prevalent type of knowledge item that the participants have ever found affected by issues was a definition, axioms, theorems, proofs, examples, notation definitions for rendering symbols, and theories also being quite common. The most common issue was that a knowledge item was simply wrong, followed by being incomprehensible, presented in an uncommon style, underspecified, redundant, or its trueness being uncertain. Issues were mostly solved by directly improving the affected knowledge item (as opposed e. g. to deleting it or creating another one), or by splitting it into more than one. Still, some participants have experienced issues being unresolved and mostly attributed this to an insufficient tool support for restructuring knowledge items.

[6] See `http://tinyurl.com/5qdetd`

Whenever there is a discourse about a knowledge item, the system will check whether there is an issue that is both unresolved (meaning that no decision on it has been posted yet) and not challenged as invalid by the existence of a majority of disagreement replies to it. If ideas have been posted on how to resolve this issue, the most popular in terms of the ratio of agreements to disagreements will be selected. Formally, any issue s satisfying $\nexists c \in \mathrm{Dec}(s) \wedge (\exists d \in \mathrm{D}(s) \Rightarrow |\mathrm{A}(s)| > |\mathrm{D}(s)|)$ is considered legitimate, and the longest discussed (i. e. oldest) idea i in $\underset{i \in \mathrm{Id}(s),\, \exists a \in \mathrm{A}(i)}{\arg\max} \frac{|\mathrm{A}(i)|}{|\mathrm{D}(i)|+1}$ wins, where Id, Dec, A, and D denote sets of ideas, decisions, agreements and disagreements with an issue or idea, respectively. The system will provide assistance to any volunteering author to implement the winning solution in the space of knowledge items, e. g. by automatically creating a template for a new knowledge item that the author can then complete. If an author follows the steps proposed by the system, the system will conclude the respective discourse by posting automatically generated decisions. Still, we leave the freedom to the community to implement solutions manually, when users feel that the automatic support is not adequate to the wickedness of the current problem. In this case, the author to resolve an issue has to document this decision manually. Any thread that has been concluded by a decision will no longer be considered by the system.

5 Implementation

We have implemented a proof of concept in the mathematical semantic wiki SWiM. SWiM extends IkeWiki (cf. section 3) with mathematical semantic markup. It supports editing, viewing, and importing/exporting mathematical knowledge that is represented in the markup languages OpenMath [7] and OMDoc[16] and has the ontologies corresponding to these languages built in. One SWiM page usually holds a mathematical statement. From a saved or imported page, SWiM extracts an RDF outline in terms of these ontologies and uses this RDF graph to supports semantic navigation through mathematical knowledge items and to offer an infrastructure for more sophisticated semantic services [19].

We have extended the user interface for discussions by the possibility to make not just untyped comments but to post issues or ideas of specific types and to state one's agreement or disagreement (fig. 2). The DILIGENT argumentation ontology, in the current prototype without the classes for fine-grained argumentation about issues and ideas, but enhanced by mathematics-specific issue and idea types (i. e. to the extent as shown in fig. 1), is preloaded into SWiM; discussion posts are represented in RDF as instances of these types. Thus, the structure of the argumentation forms an overlay network on top of the raw structure of the threads represented in SIOC. On the top level of a discussion page, the user is invited to post issues of types that are applicable to the type t_k of the knowledge item to be discussed; a reply to an issue t_{is} can have one of the idea types that are applicable to (t_k, t_{is}). Thanks to IkeWiki's built-in ontology editor, privileged members of the community can even dynamically and interactively adapt the argumentation ontology to the community's needs. The formulæ for determining unsolved legitimate issues and "winning" ideas are implemented as sequences of SPARQL [30] queries to the RDF graph of the discussion about the knowledge item currently viewed. The assistance with implementing a solution is currently hard-coded into SWiM.

6 Use case

To demonstrate the system, we consider the situation that there is an incomprehensible theorem. The user Alice wants to report that issue. She opens the discussion page for the theorem and posts a new issue. As a type of the issue, she can select any type that is applicable to theorems. Then, Bob replies by clicking the "Idea" reply button in the issue post and selecting an idea type; any type of idea that is applicable to incomprehensible theorems can be selected. Cecil disagrees with the idea, Dan agrees (fig. 2). Now assume that Eric replies to the idea with another agreement and that, after that, Frank visits the theorem: By then SWiM will have identified the idea to provide an example for the theorem as the best one to resolve the issue and display a messages that proposes this, offering a link to start a semi-automatic assistant (fig. 3). If Frank decides to provide the example and clicks on the link, a new example page, pointing to the original theorem, would be created, and he can fill out the template (fig. 3).

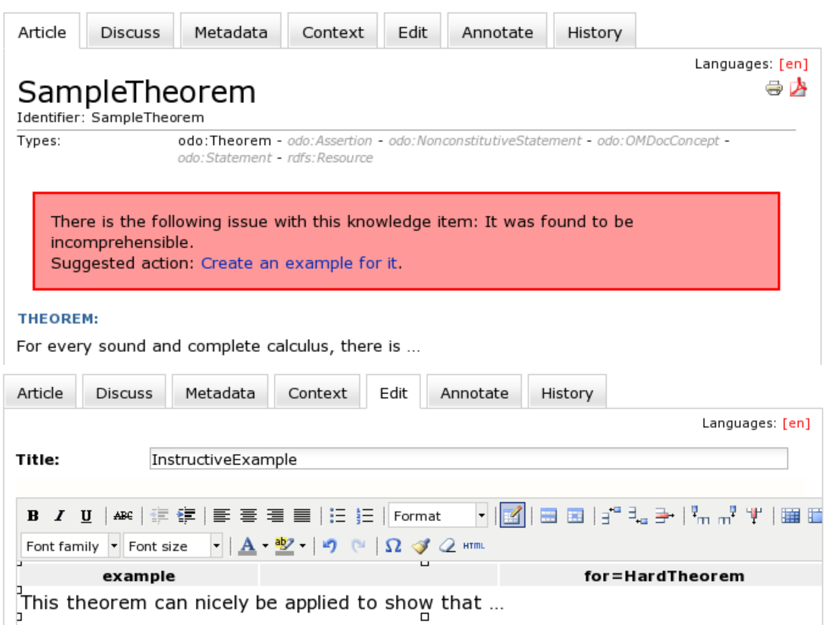

Fig. 3. Warning about an issue, and the solution offered

7 Evaluation

As this work is currently in progress, we focus our evaluation on discussing the extent of argumentation and assistance that our ontology and the implementation support. In section 9, we will provide an outlook to domain-specific case studies that we are planning to conduct ourselves. Further feedback is anticipated from users of IkeWiki, where we have recently integrated the basic argumentative discussion functionality without domain-specific extensions.

Arguments about issues and ideas have been left out in our first prototype in favour of simple agreement or disagreement, but we are planning to complete the coverage of the DILIGENT argumentation ontology. Following the findings of the DILIGENT authors, we do not, however, consider the currently restricted set of available domain-specific issue and idea types an obstacle; we think it just has to be refined to cover the most common situations in mathematical knowledge management. Concerning their restricted set of argument types (e. g. "challenge" and "justification"), the DILIGENT authors have found out in their case studies that this made discussions more effective and focused [33]; we assume that the same will turn out for our issue and idea types. As issues in our model refer to knowledge items, and the assistance the system can offer depends on the type of knowledge item, the knowledge must already be structured to some extent. There are two situations where this is likely not to be the case: When knowledge about some *new* topic has not even been conceptualised, or when it has been conceptualised and put on a wiki page, but not yet formalised (here: annotated with a type). To improve on this, we are considering to provide a global discussion space, where issues of conceptualisation can be raised, as it was originally intended with DILIGENT (cf. section 4), and to introduce a generic issue type "needs formalisation", which can be filed with any knowledge item that does not yet have a type from the domain-specific ontology. *Assistance* in the latter case would likely need techniques like natural language processing (NLP), which we have not yet considered: Once a knowledge item or discussion post has been given a type, we base all further decisions about assistance exclusively on this type. Obviously, this is only meaningful if the informal text contained in a resource does not contradict its formal type. Currently, we hold the users responsible for that, but NLP would help here, too.

8 Related work

Several projects are related to ours, each in a different way: panta rhei [27] as a different approach at a discussion platform for structured mathematical knowledge, Drupal as a content management system with a plugin for issue tracking [8], the Lekapidia case study [34] about DILIGENT argumentations in a wiki, and finally the Foucault@Wiki study [29] about discourses in Wikipedia.

panta rhei is an interactive and collaborative reader for mathematical documents, currently being evaluated in the educational context of a computer science lecture [27, 26]. The similarity with SWiM lies in the same domain of application – in fact, either system can import OMDoc – and in the possibility to create typed discussion posts about knowledge items; the key differences lie in the use of knowledge items – arranged for reading in panta rhei, whereas the main concern of SWiM is to edit them – and the different knowledge representation: SWiM relies on an underlying RDF model and an ontology for browsing, searching, and editing workflows, whereas panta rhei uses hand-made SQL database queries.

A page in panta rhei usually contains one exercise or a sequence of one to a few lecture slides, which usually contain a group of related mathematical statements, such as the definition of a few symbols together with an example where they can occur, or a theorem followed by its proof. However, content authors can assign identifiers to any subitem of a page to designate it as an annotatable knowledge item. SWiM is less flexible here: Only those knowledge items that have their own wiki page are annotatable. While imported documents are automatically split into pages of statement size – the complete document being retained as a frame that transcludes the split-off fragments for better readability –, authors would have to do this manually to achieve a finer granularity.

In panta rhei, threaded discussion items can be posted on any annotatable knowledge item; in addition, there is a global forum. Each post has to be typed as, e. g., "advice", "answer", "comment", "example", or "question"[7]. The set of possible types of a post is not restricted by the type of the knowledge item or the post it replies to. Compared to SWiM's use of an argumentation ontology, which encourages a targeted discussion towards solutions, this potentially makes discussion threads less focused. panta rhei currently uses the types of posts for statistical purposes and for search. Statistics are currently not computed in SWiM; semantic search is possible and powerful thanks to the support for inline SPARQL queries, but not yet friendly to users who do not know SPARQL. Another annotation-related feature of panta rhei without a SWiM counterpart is the ability to *rate* knowledge items on a scale from 0 to 10 w. r. t. several measures like their difficulty or helpfulness. While the argumentations in SWiM implicitly rate knowledge items, there is no agile, one-click user interface for this.

Theoretical work has been done towards supporting mathematical communities of practice (CoP) with panta rhei [25]. One goal is to identify subcommunities among the users according to the documents they rate as relevant for themselves, based on an extensional model of CoP as document collections as proposed in [15]. SWiM currently follows the simple-minded approach of considering the whole user base of a site as one CoP, not differentiating between the activities of different users.

Drupal is a modular framework and content management system [8]. The main strength of Drupal is its flexibility through a great variety of modules which allows it to be able to cater for almost any need. Drupal's eco-system is mainly composed of nodes, comments and taxonomy vocabularies. Each of these can easily be enhanced to behave in various ways. The primary type of content is

[7] This is a custom vocabulary made for panta rhei, not an argumentation ontology.

a node. Out of the box, Drupal can let users comment on the content of the node, or reply to other comments resulting in threaded conversations – which is unsupported by most wikis. Drupal also offers many hooks to let developers jump in during the page generation, validation or submission process and that way react precisely according to the input of the user. The project issue tracking module is already using these features, along with some categorisation[8].

Lekapidia was a case study conducted by the authors of the DILIGENT argumentation ontology. They preloaded a semantic wiki (coefficientMakna) with the argumentation ontology and replayed the collaborative engineering of a simple dessert recipe ontology, which had earlier been developed using the DILIGENT methodology, in their wiki, and found out that the wiki "significantly reduces the effort to capture the arguments in a structured way" [34]. A collection of semantically structured mathematical knowledge can also be considered an ontology, particularly if it contains definitions, i. e. introduces new mathematical concepts in a formal way. The argumentation ontology in SWiM is particularly enhanced for the domain of mathematics, which allows for creating more specific issues and ideas in a machine-understandable way and for implementing semi-automatic assistance in implementing certain of these specific ideas into the system.

Foucault@Wiki was a study that analysed discourses in Wikipedia: how editing summaries and discussion posts relate to changes made to Wikipedia articles, and what types of changes occurred [29]. From this, a model for argumentations in Wikipedia, in fact an informal site-specific argumentation ontology, was derived. However, this analysis had to be made without machine support, as Wikipedia articles are largely unstructured, and discussions and editing summaries are given in natural language and the space of possible arguments is unrestricted, as opposed to a finite set of types in a formal argumentation ontology. Note that the goal of this study was not to design software support for discourses or for improving knowledge items, i. e. Wikipedia articles. We consider such a goal extremely hard to achieve, as discussion pages in the MediaWiki engine are not structured and articles frequently contain multiple section that do not focus on one knowledge item only, which makes it hard to express what part of an article a discussion post refers to.

9 Conclusion and Outlook

We have justified the need for structured and annotated discussions in a community-driven knowledge base. With a domain-specific argumentation ontology, we have covered common types of issues that can occur with mathematical knowledge items, common types of ideas on how to solve these issues, and a semi-automatic assistance in implementing a solution approved by the community. We have implemented and demonstrated a first proof of concept in the OMDoc-based semantic wiki SWiM but consider the methodology easily transferable to other

[8] See `http://drupal.org/project/project_issue`

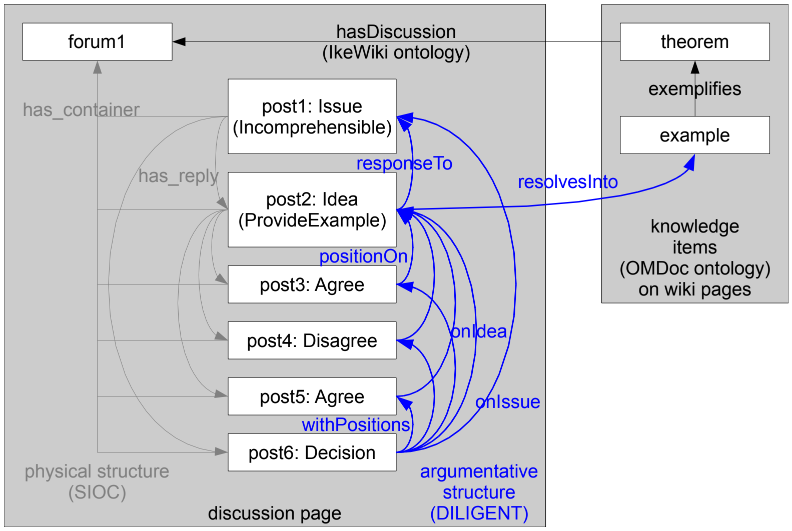

Fig. 4. RDF graph of a sample use case

community tools, as it is largely based on an abstract ontology. SWiM itself assumes that all users of a site form one community of practice (CoP [37]). It strengthens this CoP through an improved documentation of practices: The discourses that led to creation or revision of knowledge items are recorded in a structured way and can be traced back transparently from the affected knowledge items, thus allowing new members to learn more about established practices within the CoP. This aspect of "practice as learning" [37] has been confirmed as essential for mathematical CoP [15].

As our system is built on RDF, Semantic Web technologies can be used in the implementation as well as in integration with other systems. For example, the Exhibit framework [13] can consume RDF data and provide us with a timeline visualisation of the argumentation process (fig. 5). If we follow the Linked Data guidelines for Semantic Web publishing, the knowledge items and discussions in our system will be visible to Linked Data crawlers such as Sindice [35] and browsers such as Tabulator [2]. If we can crosslink with other sites, discussions can refer to knowledge items across system boundaries.

Two practical settings, where we will evaluate our system, are the Flyspeck and OpenMath projects. Flyspeck is a large-scale proof formalisation effort concerned with developing a machine-verifiable representation of a proof of the Kepler sphere packing conjecture [22]. We have gathered first requirements for supporting this project with a wiki and consider support for discussing formalisation issues and for refactoring formal structures highly important. In OpenMath, SWiM will be used for editing content dictionaries (collections of definitions of mathematical symbols). There, conceptualisation and formalisation of symbols and their notations needs to be supported [20]. A key result that we

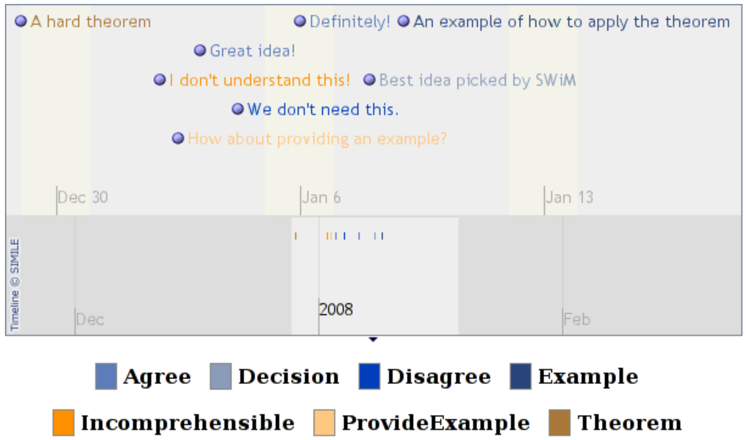

Fig. 5. The Exhibit timeline view of the argumentation in fig. 4, starting at a hard theorem and ending at a helpful example.

expect from these case studies is to what extent the automated identification of "winning" solutions and the support offered for implementing them will be found satisfactory by knowledge engineers – whether the 80/20 rule will apply (i. e. 80% of the everyday issues will be solvable with semi-automatic assistance implemented for 20% of all possible solutions), or whether a large part of issues turns out to be too wicked. In the latter case we still hope that the argumentation ontology will support a more focused and productive discussion about wicked problems than in a system with unstructured discussions and thus facilitate finding and implementing solutions even without further automatic assistance.

Acknowledgments The authors would like to thank John Breslin for advice with SIOC and issue-based argumentation, Markus Luczak-Rösch for guidance with using the DILIGENT argumentation ontology, Christine Müller for helping to investigate similarities and differences of our system and panta rhei, Max Völkel for ideas how to make use of argumentation in a wiki, and Uldis Bojārs for suggestions how to extend our approach to the web.

References

1. J. Baumeister, J. Reutelshöfer, K. Nadrowski, and A. Misok. Using knowledge wikis to support scientific communities. In Müller [24].
2. T. Berners-Lee, Y. Chen, L. Chilton, D. Connolly, R. Dhanaraj, J. Hollenbach, A. Lerer, and D. Sheets. Tabulator: Exploring and Analyzing linked data on the Semantic Web. In *Proceedings of the The 3rd International Semantic Web User Interaction Workshop (SWUI06)*, Nov 2006.
3. U. Bojārs and J. G. Breslin. SIOC core ontology specification. W3C Member Submission, World Wide Web Consortium (W3C), June 2007. `http://www.w3.org/Submission/2007/SUBM-sioc-spec-20070612/`.

4. A. Brändle. Zu wenige Köche verderben den Brei. Eine Inhaltsanalyse der Wikipedia aus Perspektive der journalistischen Qualität, des Netzeffekts und der Ökonomie der Aufmerksamkeit. Master's thesis, Universität Zürich, 2005.
5. J. G. Breslin, S. Decker, A. Harth, and U. Bojārs. SIOC: an approach to connect web-based communities. *International Journal of Web Based Communities*, 2(2):133–142, July 2006.
6. M. Buffa, F. Gandon, G. Ereteo, P. Sander, and C. Faron. Sweetwiki: A semantic wiki. *Web Semantics: Science, Services and Agents on the World Wide Web*, 2008.
7. S. Buswell, O. Caprotti, D. P. Carlisle, M. C. Dewar, M. Gaetano, and M. Kohlhase. The Open Math standard, version 2.0. Technical report, The Open Math Society, 2004. `http://www.openmath.org/standard/om20`.
8. Drupal.org – community plumbing. web page at `http://drupal.org`.
9. A. Ebersbach, M. Glaser, and R. Heigl. *Wiki: Web Collaboration*. Springer-Verlag New York, 2008.
10. D. Fensel. *Ontologies: A Silver Bullet for Knowledge Management and Electronic Commerce*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2003.
11. Friend of a Friend (FOAF) project. `http://www.foaf-project.org/`, seen June 2008.
12. T. R. Gruber. A translation approach to portable ontology specifications. *Knowledge Acquisition*, 5(2):199–220, June 1993.
13. D. Huynh, D. Karger, and R. Miller. Exhibit: Lightweight structured data publishing. In *16th International World Wide Web Conference*, Banff, Alberta, Canada, 2007. ACM.
14. A. Kittur, B. Suh, B. A. Pendleton, and E. H. Chi. He says, she says: conflict and coordination in wikipedia. In M. B. Rosson and D. J. Gilmore, editors, *CHI*, pages 453–462. ACM, 2007.
15. A. Kohlhase and M. Kohlhase. Communities of Practice in MKM: An Extensional Model. In J. Borwein and W. M. Farmer, editors, *Mathematical Knowledge Management, MKM'06*, number 4108 in LNAI, pages 179–193. Springer Verlag, 2006.
16. M. Kohlhase. OMDoc – *An open markup format for mathematical documents [Version 1.2]*. Number 4180 in LNAI. Springer Verlag, 2006.
17. W. Kunz and H. W. J. Rittel. Issues as elements of information systems. Working paper 131, Institute of Urban and Regional Development, University of California, Berkeley, July 1970.
18. C. Lange. The OMDoc document ontology. web page at `http://kwarc.info/projects/docOnto/omdoc.html`, seen November 2007.
19. C. Lange. Towards scientific collaboration in a semantic wiki. In A. Hotho and B. Hoser, editors, *Bridging the Gap between Semantic Web and Web 2.0 (SemNet 2007)*, 2007.
20. C. Lange. Mathematical Semantic Markup in a Wiki: The Roles of Symbols and Notations. In Lange et al. [23].
21. C. Lange. SWiM – a semantic wiki for mathematical knowledge management. In S. Bechhofer, M. Hauswirth, J. Hoffmann, and M. Koubarakis, editors, *ESWC*, volume 5021 of *Lecture Notes in Computer Science*, pages 832–837. Springer, 2008.
22. C. Lange, S. McLaughlin, and F. Rabe. Flyspeck in a semantic wiki – collaborating on a large scale formalization of the Kepler conjecture, June 2008.
23. C. Lange, S. Schaffert, H. Skaf-Molli, and M. Völkel, editors. *3rd Workshop on Semantic Wikis*, Costa Adeje, Tenerife, Spain, June 2008.
24. C. Müller, editor. *1st Workshop on Scientific COmmunities Of Practice (SCooP-2007)*, 2007.

25. C. Müller. Towards the identification and support of scientific communities of practice. In *SCooP 2007* [24].
26. C. Müller. The CS precourse project. http://cs-precourse.kwarc.info/, seen June 2008.
27. C. Müller and M. Kohlhase. panta rhei. In A. Hinneburg, editor, *Wissens- und Erfahrungsmanagement LWA (Lernen, Wissensentdeckung und Adaptivität) conference proceedings*, pages 318–323, 2007.
28. E. Oren, R. Delbru, K. Möller, M. Völkel, and S. Handschuh. Annotation and navigation in semantic wikis. In M. Völkel, S. Schaffert, and S. Decker, editors, *Proceedings of the 1st Workshop on Semantic Wikis, European Semantic Web Conference 2006*, volume 206 of *CEUR Workshop Proceedings*, Budva, Montenegro, June 2006.
29. C. Pentzold and S. Seidenglanz. Foucault@Wiki – first steps towards a conceptual framework for the analysis of wiki discourses. In D. Riehle and J. Noble, editors, *Proceedings of the 2006 International Symposium on Wikis*, ACM Press, Aug. 2006.
30. E. Prud'hommeaux and A. Seaborne. SPARQL query language for RDF. W3C Recommendation, World Wide Web Consortium, Jan. 2008. Available at http://www.w3.org/TR/2008/REC-rdf-sparql-query-20080115/.
31. H. W. J. Rittel and M. M. Webber. Dilemmas in a general theory of planning. *Policy Sciences*, 4(2):155–169, June 1973.
32. S. Schaffert. IkeWiki: A semantic wiki for collaborative knowledge management. In *1st International Workshop on Semantic Technologies in Collaborative Applications STICA 06, Manchester, UK*, June 2006.
33. C. Tempich, H. S. Pinto, Y. Sure, and S. Staab. An argumentation ontology for DIstributed, Loosely-controlled and evolvInG Engineering processes of oNTologies (DILIGENT). In A. Gómez-Pérez and J. Euzenat, editors, *ESWC*, volume 3532 of *Lecture Notes in Computer Science*, pages 241–256. Springer, 2005.
34. C. Tempich, E. Simperl, M. Luczak, R. Studer, and H. S. Pinto. Argumentation-based ontology engineering. *IEEE Intelligent Systems*, 22(6):52–59, 2007.
35. G. Tummarello, R. Delbru, and E. Oren. Sindice.com: Weaving the open linked data. In K. Aberer, K.-S. Choi, N. F. Noy, D. Allemang, K.-I. Lee, L. J. B. Nixon, J. Golbeck, P. Mika, D. Maynard, R. Mizoguchi, G. Schreiber, and P. Cudré-Mauroux, editors, *ISWC/ASWC*, volume 4825 of *Lecture Notes in Computer Science*, pages 552–565. Springer, 2007.
36. Resource Description Framework (RDF). http://www.w3.org/RDF/, 2004.
37. E. Wenger. *Communities of Practice: Learning, Meaning, and Identity.* Cambridge University Press, 2005.
38. Wikipedia, the free encyclopedia. http://www.wikipedia.org, 2001–2007.
39. Wikipedia: Bewertungsbausteine (aus Wikipedia, der freien Enzyklopädie). http://de.wikipedia.org/w/index.php?title=Wikipedia:Bewertungsbausteine%&oldid=44205795, Mar. 2008.
40. Wikipedia: Diskussionsseiten (aus Wikipedia, der freien Enzyklopädie). http://de.wikipedia.org/w/index.php?title=Wikipedia:Diskussionsseiten&o%ldid=46659885, May 2008.
41. Wikipedia: Neutraler Standpunkt (aus Wikipedia, der freien Enzyklopädie). http://de.wikipedia.org/w/index.php?title=Wikipedia:Neutraler_Standpunk%t&oldid=46873335, June 2008.
42. Hilfe: Zusammenfassung und Quelle (aus Wikipedia, der freien Enzyklopädie). http://de.wikipedia.org/w/index.php?title=Hilfe:Zusammenfassung_und_Que%lle&oldid=46731506, May 2008.