

Collaborative Information Finding in Smaller Communities: The Case of Research Talks

Peter Brusilovsky, *Member, IEEE*, Denis Parra, *Member, IEEE*, Shaghayegh Sahebi, and Chirayu Wongchokprasitti

Abstract—Social navigation and social tagging technologies enable user communities to assemble the collective wisdom, and use it to help community members in finding the right information. However, it takes a significantly-sized community to make a social system truly useful. The question addressed in this paper is whether collaborative information finding is feasible in the context of smaller communities. To answer this question, we developed two social systems specifically focused on smaller communities – CoMeT and Conference Navigator II – and explored several techniques to increase the volume of user contributions. This paper reviews the explored techniques and presents empirical evidence that demonstrate their effectiveness.

Index Terms— Keywords, Papers, Recommendation System, Talks, Tags, User Profile Fusion

I. INTRODUCTION

The modern social Web demonstrates that a community of users possesses the *collective wisdom* that can help its members to deal with information overflow. Through social bookmarking and tagging, the users implicitly help each other by contributing interesting and useful information and categorizing it to make it easier to find. Existing social tagging systems achieved great success in helping various communities to collect, organize, and share such items as bookmarks (Delicious), news items (Digg), pictures (Flickr), and research papers (CiteULike).

While different from each other, the social tagging systems mentioned above share a very important feature: their success is based on the volume of contributions from their users. The users have to post and tag items to make these systems effective. Unfortunately, experience with many online communities shows that a relatively small fraction of their members become active contributors [1]. As a result, it takes a community of a significant size to make a collaborative tagging system truly useful. Not surprisingly, the successful systems mentioned above feature thousands to millions of users from many countries.

Manuscript received August 20, 2010. This work was partially supported by National Science Foundation under Grant IIS-0447083.

P. Brusilovsky, Denis Parra, and Chirayu Wongchokprasitti are with the School of Information Sciences and the Intelligent Systems Program, University of Pittsburgh, Pittsburgh, PA 15260 (412-624-9404 fax: 412-624-2788; e-mail: peterb@pitt.edu; dap89@pitt.edu, chw20@pitt.edu).

S. Sahebi is with the Intelligent Systems Program, University of Pittsburgh, Pittsburgh, PA 15260 (e-mail: sahebi@cs.pitt.edu).

Naturally, small communities that range from several dozens to several hundreds users could also greatly benefit from the collaborative information finding power of social tagging systems: however, is it really feasible? Is it possible to engineer a social tagging system that can be powered by a non-traditionally small number of users? For the last five years, we have been trying to answer this research question for a special domain of research talks. We developed two social systems for collaborative sharing and finding of *research talks*. One of these systems can work for the communities of this size. The CoMeT system targets researchers at the University of Pittsburgh and Carnegie Mellon University campuses enabling them to share information about colloquia, thesis defenses, and other kind of talks regularly held at such research universities. The Conference Navigator II (CNII) system targets researchers attending a single academic conference helping them to discover and share interesting papers, posters, and workshop talks presented during the conference. For both systems, the size of the target community is very small by the standards of social tagging systems. To make these systems more sustainable, we focused our efforts on increasing the volume of social contributions coming from the target communities. We explored several ideas and system design features targeting two goals: encouraging a larger proportion of users to become contributors and collecting more contributions from individual community members. This paper reviews some of the explored ideas and reports empirical evidence collected by both systems. Following a brief review of similar research, the paper presents the story of our work in three parts. The “past” section presents our core design and offers some empirical evidence in its favor. The “present” section reports the results of our most recent work on adding personalized recommendations with the resulting empirical evidence. The “future” section presents the “early recommendation” approach, which we plan to explore in the near future and offers the results of the predictive evaluation of this approach. We conclude with a brief summary of our work and the challenges to be addressed in future research efforts.

II. INCREASING USER PARTICIPATION IN SOCIAL SYSTEMS

It has been recognized that success of all kinds of social software and online communities is highly dependent on participation of their users. This recognition brought significant interest to the problem of user participation and

increased the volume of research on this problem. One of the most frequently cited issues in the area of user participation is the under-contribution and inequality of contribution to online communities. In most online communities, 1% of users account for 90% of content [2]. In Wikipedia, one of the most successful examples of social systems, only 0.01% of users are active contributors¹. While bigger communities like Wikipedia with larger number of users can survive even with a small percentage of users contributing, the survival of small online communities with a small number of users is highly dependent on contributions from the majority of the users. However, even large communities can be affected by the under-contribution problem through participation inequality bias (a small percentage of the population represents the views of the entire population). Not surprisingly, the research on user participation and the approaches to increase it have become a hot topic for both social science and computer science researchers [1].

As summarized in an earlier paper [3], community contribution can be motivated by extrinsic and intrinsic factors. Currently, the majority of research efforts focus on the exploration of extrinsic factors such as rewards [4-8], reputation [9, 10], competitive games [11, 12], and community value [1]. In addition, most of the research on community contribution has been done in the context of large communities. Our work differs from the mainstream in both aspects. It targets smaller communities and focuses mostly on intrinsic factors, especially the commonly referenced as so-called “personal need” factor [13, 14]. The idea of the “personal need” approach is to turn an activity useful for the whole community into an activity that is meaningful for the user. For example, in our work on our course recommendation system CourseAgent [13], we explored the use of personal course planning to make the course rating activity beneficial to the user personally, not just to the community.

III. THE PAST: THE DESIGN OF THE SYSTEMS

The work on the Conference Navigator II and the CoMeT systems has been motivated by our earlier experience with “the original” Conference Navigator system [15] in 2005-2007. The original Conference Navigator was designed as a social navigation system, which draws the attention of its users to highly-ranked papers and sessions at a conference. As an implicit sign of interest [16] the system used traditional “hit rate”. In addition, the system pioneered the use of user “scheduling activity” as an explicit sign of interest. It allowed users to create their own personal conference schedule by adding paper presentations that they would like to attend and making decisions in case interesting papers are presented at the same time. Scheduling is a good example of the “personal need” approach. Users add papers and sessions to their schedules to meet their own needs – a schedule created in advance is a great help in navigating a busy multi-stream conference. Yet, the fact that the paper is added to someone’s

schedule provides an important piece of community wisdom and allows the system as a whole to identify more interesting papers.

Both CNII and CoMeT attempted to expand upon this original Conference Navigator work by merging the social navigation features of this system with features typical for social tagging systems and exploring the “personal need” approach in the new context.

A. The Design of the Conference Navigator System

Conference Navigator II (CNII) is a system developed to assist participants at an academic conference. The goal of the system is to help its users to find the most interesting and relevant talks. CNII is accessed through a Web-based interface that provides information about the conference schedule along with additional details about individual papers. The users can browse through this information in several ways. The social navigation feature of the system guides its users to papers, which have been identified as most interesting for the whole set of system users as well as for the specific community. By community, we mean a relatively small group of users with common interests with respect to the domain of the conference[17].

As a system that blends social tagging with social navigation, CNII uses several types of interest indicators to drive the social navigation functionality. In addition to implicit interests derived from user browsing activity (details about which papers were opened), it tracks three other kinds of user activities which the system encourages the user to do: scheduling (adding a paper to personal schedule), tagging, and contributing to a community. The latter action means that the user explicitly adds a paper to one of the existing communities and can be considered as a blend of group-level social navigation as used in [14, 18] and regular tagging.

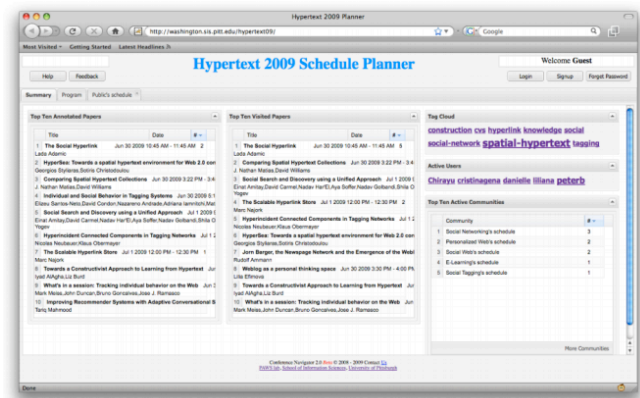


Fig. 1. The Conference Summary Tab in CNII.

The system offers a tab-based interface, which can feature several types of tabs. The users start with the Conference Summary tab (Fig. 1). It shows a summary of the conference: top-ten annotated papers, top-ten visited papers, a tag cloud, an active users cloud, and top-ten active communities. The Program Browser tab provides detailed information about conferences sessions and papers. The Schedule Tab can display the personal schedule of the user, other users’

¹ <http://en.wikipedia.org/wiki/Wikipedia:About>

schedules, or the schedule of one of the system's communities.

Each of these tabs provides different ways to access individual papers. Clicking on the paper title opens the Paper Panel with details about the paper (title, authors, abstract) and associated social information (users' tags associated with the paper, users who bookmarked the paper and communities to which the paper has been added). This is also the place where users can provide all kinds of contributions: schedule the paper, add notes and tags, and recommend it to one or more communities (Fig. 2).

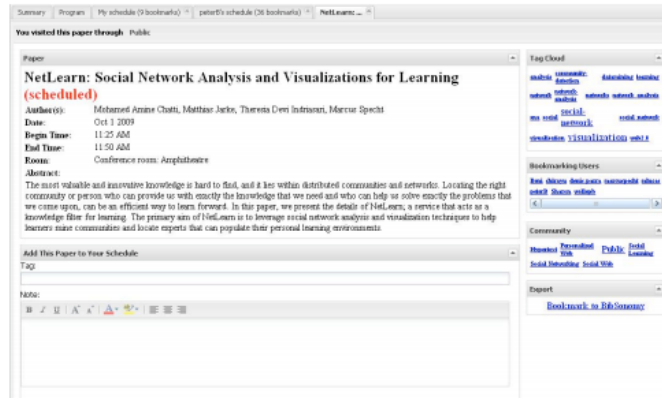


Fig. 2. Paper panel in CNII with interface for user contribution.

B. The Design of the COMET System

CoMeT is a collaborative system for sharing information about research colloquia. The system targets researchers and students at the University of Pittsburgh and Carnegie Mellon University, which are located next to each other in the city of Pittsburgh, Pennsylvania. It supports both passive and active dissemination. Every user can post a talk by submitting a simple form. Mandatory fields include title, speaker, date, time, and location. The talk description field is not mandatory, although almost all posted talks included an abstract and information about the speaker in the description. Users can browse posted talks by date (Fig. 3), standing series, organizing departments, as well as by other ways. The system can also disseminate talks using iCal, Google Calendar and RSS feeds.

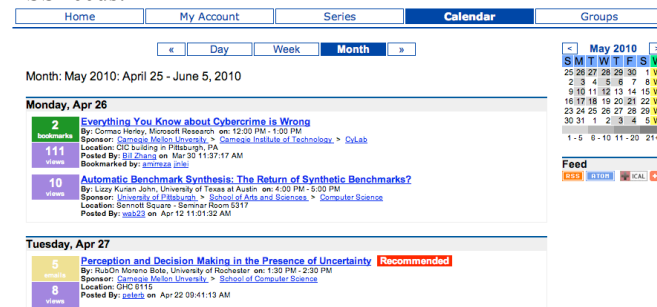


Fig. 3. A Calendar view of CoMeT interface.

When users find interesting talks, they can add them to their schedule, add tags and comments, and recommend them to one or more communities (Fig. 4). These features are directly modeled after CNII. An additional feature of CoMeT is the ability to share information about talks with their friends by

email. Also, to enhance the social navigation features of CoMeT, the cumulative user activity related to a talk (viewing, tagging, or sending by e-mail) is visualized in all lists where the talk is displayed (Fig. 4).




Fig. 4. Bookmark the Talk.

C. The Evaluation of the Design: CNII

The design of CNII has been finalized after its pilot at Adaptive Hypermedia 2008. After that, the system has been used at three conferences in 2009: User Modeling, Adaptation, and Personalization 2009 (UMAP 2009), Hypertext 2009 (HT 2009), and European Conference on Technology Enhanced Learning 2009 (ECTEL 2009). The enhanced version of the system has also been used at UMAP 2010.

We used the experience gained at these conferences to perform the evaluation of our core design features. Our main concern was the volume of user contributions. Was the “personal need” approach based on scheduling sufficient to collect an effective level of social feedback? How large was the volume of tagging and community contributing, since these activities were not specifically motivated?




Fig. 5. The graph shows the number of registered users, active users, and then users engaged in different contributory activities. Numbers in parenthesis represent the number of attendees at each conference.

A frequent yardstick for measuring participation inequality [1] is Nielsen's 90-9-1 rule [2]: 90% of users are lurkers, 9% contribute a small proportion, and 1% contributes the most content. As CNII data shows, at all conferences, the system managed to achieve a much higher level of user contributions, compared to Nielsen's model, as can be seen in Fig. 5. In this figure, registered users are all those who have created an account in CNII and active users are those who have undertaken some activity in addition to simple browsing: schedule a paper, tag a paper or contribute papers to communities. As the data shows, at every conference, the

number of *active* users – and, in particular, scheduling users – was over 50% of the *registered* users. We can consider it to be evidence of the success of our efforts to engage users in the scheduling process through the “personal need” approach. However, was it really the personal scheduling need that increased user engagement?

Fig. 6 provides another view of this data showing only three 2009 conferences at which we utilized the same version of the system and measuring the number of users engaged in each kind of contributory activity as a percentage of the entire number of registered users. This figure reveals a clear difference between ECTEL 2009 (almost all registered users scheduled at least one talk), HT 2009 (around 96% of users scheduled at least one talk), and UMAP 2009 (just over 50% of users scheduled at least one talk). We attribute this result to the fact that there are differences in the number of parallel streams in each conference: HT 2009 was a 2-stream conference and ECTEL 2009 was a 3-stream conference, as compared to the 1-stream UMAP 2009. The multiple streams made the process of creating a personal schedule much more important in order to manage overlapping talks. This data provides additional evidence that the volume of scheduling activity was driven by the user’s personal scheduling need.




Fig. 6. Percentage of scheduling users, tagging users, and community contributors. Numbers in parenthesis represent the number of registered users and the number of attendees at the conference.

We observed that user engagement in tagging and community contribution was always less than the number of scheduling users. It can be explained by the fact that our 2009 design did not provide incentives to engage users in these activities. While scheduling was important to efficiently navigating the conference, the usefulness of the other two activities was not obvious to users. We can also see that the proportion of more active contributing users to scheduling users is relatively stable and is not affected by the conference size (Fig. 7). We believe that this is also a result of our 2009 design: in order to tag or contribute to a community, a user had to pull up the scheduling interface, where tags and communities were options.




Fig. 7. Percentage of scheduling users that also tagged (blue) or contributed to communities (red). Numbers in parenthesis represent the number of scheduling users and the number of attendees at the conference.

D. The Evaluation of the Design: CoMeT

CoMeT system was launched in August 2009 and is still in operation. Since most of the research talks happen during the fall and spring semesters, the majority of the talks were posted on CoMeT between August 2009 and May 2010. At the moment, CoMeT has 209 registered users, but only 35 can be considered to be active contributors. To be considered as active, users must have performed some direct contribution activity such as bookmarking a talk, sending a talk by e-mail, commenting on a talk, contributing it to a community, or rating it. The design of the system ensured that even users who only browsed the talks become indirect contributors; however, just viewing a talk is not considered to be an active contribution.

While only users with registered accounts can do the aforementioned activities for a talk, all users including guest users can view all of the colloquia. As a result, we don’t have the exact number of system visitors including guests or silent users. In Table I, which counts the different IP addresses of visiting users, we can see that users (both registered and guest) have checked the schedule of talks from 5355 different IP addresses and viewed details of at least one talk from 1959 different IP addresses. CoMeT’s 209 registered users have logged in to the system via 345 distinct IP addresses. The fraction of users to IP addresses implies a much larger number of non-registered “lurkers” using the system. Altogether, we can see a remarkably high activity level.

TABLE I
NUMBER OF DISTINCT IP ADDRESSES CONNECTING TO CoMeT

	Number of Distinct IP Addresses
Visiting the System	5355
Viewing a Talk	1959
Logging in to the System	345

We can see the distribution of different activities and colloquia viewing among registered users in Table II and Fig. 8. They both show that 43.5% of registered users have viewed the details about some of the talks, while 16.7% of them provided a contribution, based on our definition. Most of the active users were rating and/or bookmarking talks; very few contributed by commenting on the talks.

TABLE II
NUMBER AND PERCENTAGE OF REGISTERED USERS OF CoMeT BASED ON DIFFERENT ACTIVITIES

	Number of Active Users	Percentage of Registered Users
Having Account Users	209	100
Viewing Users	91	43.54
Active Users	35	16.75
Bookmarking Users	28	13.4
Tagging Users	21	10.05
E-mailing Users	13	6.22
Comm. Contributors	18	8.61
Rating Users	28	13.4
Commenting Users	5	2.4




Fig. 8. Number of registered users of CoMeT based on different activities

It is also possible to measure the user activity by the number of bookmarked, tagged, or community-contributed talks. As we can see in Table III (columns *a* and *b*) and in Fig. 9, all the talks (colloquia) in the system have been viewed. The most frequent activity was bookmarking a talk; it was done for 28.4% of the talks. Rating was done for 25.92% of the talks, being the second most popular activity on the talks. With just a 1.84% of talks, commenting was the least frequent activity. Also, Table III column C shows that “visiting users” (i.e., those who visited at least one talk when logged in) have visited 7.76 talks, bookmarked 2.2 talks and rated 2.01 talks in average.

TABLE III
A AND B: NUMBER AND PERCENTAGE OF ACTIVITIES ON TALKS OF CoMeT; C: AVERAGE ACTIVITY ON TALKS PER USER

	a) Talks with Activities	b) Percentage of Talks with Activities	c) Average Activity per Visiting User
Total Talks	706	100%	-
Viewed Talks	706	100%	7.76
Bookmarking Talks	201	28.47%	2.20
Tagged Talks	128	18.13%	1.40
Comm. Contributed Talks	131	18.55%	1.44
Rated Talks	183	25.92%	2.01
E-mailed Talks	133	18.84%	1.46
Commented Talks	13	1.84%	0.14




Fig. 9. Number of activities on talks of CoMeT

Comparing Table II and Table III, we can see that while the number of e-mailed talks is equal to the number of tagged talks, the number of tagging users is more than the number of e-mailing users. It shows that the number of e-mailing users is less than tagging users but these users are more active. The same is true for community contributions. The number of community-contributed papers is the same as the number of tagged papers. Yet, the number of community contributing users is less than the number of tagging users, i.e., contributing users are more active. This data allows distinguishing tagging from community contributing and e-mailing. The data hints that tagging is a regular activity, which many users do for themselves. In contrast, e-mailing and community recommendation are activities done “for others”. While fewer users are thinking about others, they are also more active. This result show that different approaches have to be used to increase the number of these contribution. To further increase tagging, we should use “personal value” approach making this activity more useful for user’s own needs, while to increase the volume of e-mails and recommendations, we should use a variety of “community value” approaches.

IV. THE PRESENT: ADDING RECOMMENDATIONS

As we found during previous evaluations, using scheduling as a personal incentive results in a relatively large volume of contributions. Yet, the volume of contributions still was not sufficient. Many users – both registered and not registered – were using the systems simply as a source of information without contributing to the community’s wisdom. It was especially evident in such contexts as the single-stream UMAP 2009 conference, where there was no overlap between talks (a single stream) and so the benefit of making personal schedules was relatively low, and the CoMeT system, where the density of talks was much lower than during any conference. It was apparent that some additional motivation is required to encourage users to contribute. Following our “personal need” approach, we decided to explore personal recommendations as a way to encourage users to do more contributing. The idea was to build in recommender functionality, which uses information about user’s currently scheduled talks in order to build a profile of interests and to recommend other talks relevant to user interests. In this situation, the users are encouraged to contribute in order to receive any recommendations and motivated to contribute

more to receive better recommendations. To explore this idea, we developed a new version of CNII, which can deliver recommendations. In addition to this major change, the new version separated the interfaces for tagging, scheduling and contributing to communities. Now, it is possible to tag or contribute a talk to a community without adding it to a personal schedule. It was done to explore the “true engagement” of users into these activities. The remaining part of this section presents the design of our recommendation mechanism and reports the results of testing the enhanced version at UMAP 2010 conference. Note that the same recommendation approach was developed for CoMeT, but we have not had a chance to evaluate it yet due to the summer 2010 break.

A. The recommendation approach in CoMeT and CNII

The current recommendation approach used in CoMeT and CNII is a simple content-based method using the abstract and title of each talk or paper (documents). To represent documents, we use keywords extracted from documents’ titles and abstracts. Each document is considered as a bag of words and represented as a vector in the keywords vector space weighted by the TF.IDF weighting scheme ($d_c = (w_{1,c}, w_{2,c}, \dots, w_{l,c})$ where $w_{i,c}$ shows the weight of i^{th} keyword in document c and l is the total number of keywords).

All documents are then represented in an $e \times m$ matrix D_t (with e as the total number of documents and m as the total number of keywords).

To recommend related documents to users, we have used the *K-nearest neighbor* method. In this method, the top K documents closest to the user profile are recommended to each user. To do this, we have to represent user profiles in the documents’ vector space.

User profiles are built based on the content of users’ bookmarked, scheduled and rated documents. We represent each user’s bookmarked and rated documents, weighted by user ratings, in a vector in the documents’ vector space. To obtain keyword-based user profiles, we used the aforementioned document representation model. The Keyword-based user profile (UP) is obtained by multiplying the vector of user documents in the documents’ vector space (U) by the matrix of document keywords represented in keywords and tags vector space (D):

$$UP = U.D' \quad (1)$$

The resulting user profile (UP) is a vector consisting of user’s related keywords, weighted based on the importance of each keyword. In this case, where user profiles are represented in the same vector space as system’s documents, we can find K closest talks/papers, which the user has not seen yet, and recommend them to that user. To measure the distance between documents and user profiles, we use the *cosine distance measure*.

B. The Effect of Recommendations: The Log Analysis

A comparison of the logs of the two versions of CNII used at UMAP 2009 and UMAP 2010 provides a good opportunity to evaluate the impact of recommendation and other new

features built into the UMAP 2010 version. UMAP 2009 and 2010 are as similar as two conferences can be with the same cohort of users, of papers, and organization. As the data shows, UMAP 2010 had fewer papers and attendees (112 and 130 vs. 133 and 170) than the 2009 conference, yet the proportion of active contributors and the volume of participation increased remarkably (Fig. 5). This increment is more visible in relative numbers, as a fraction to the total number of the attendees. As shown in Fig. 10, the new version of CNII attracted almost 30% of the users registered at UMAP 2010 compared to just 18% of at UMAP 2009. A larger percentage of the users were active in UMAP 2010 than in 2009, and most important, 22% of all attendees were scheduling papers in spite of the fact that this was a 1-stream conference. The new version saw a doubling in the number of active contributors.




Fig. 10. Percentage of papers viewed, scheduled, tagged and contributed to communities in UMAP 2009 and UMAP 2010. Numbers in parenthesis represent the total number of papers at the conference.




Fig. 11. The total number of contributory activities. Number of attendees and papers for the conferences are shown in brackets. Note that although some of co-authors of this paper attended UMAP 2009 and 2010 -and used the system “for real”- their activities were excluded to avoid “authors’ bias”.

As the fraction of actively contributing users increased remarkably, we might expect the growth in total number of contributions. As shown on Fig. 11, however, only the number of scheduling activities surged. This increase is even more impressive in relative numbers both per attendee and per paper as UMAP 2010 had both fewer attendees and fewer papers (i.e. the use of the “recommendation incentive” almost doubled the amount of collected “scheduling” wisdom per paper). In contrast, the number of tagging actions and community contributions stayed about the same in relative numbers (decreasing in absolute numbers as shown in Fig 11) despite the increase in the number of active users. We think

that it is a result of the separation of scheduling, tagging, and community recommendation interfaces. As already explained, in the 2009 version of CNII, tagging and recommendation were combined in the scheduling interface. Apparently, this layout pushed more users to complete these parts of the form. In contrast, tagging and contribution as independent activities were not as valuable for the users (neither was it motivated by the incentives). As we see in Fig 4, tagging and community contribution were of little value to users. For example, there were users who did not tag anything, yet added tags and suggested papers to communities. Yet, it has not compensated for the lost of activity. In light of this data, in the future version, we will offer both options: tagging and community recommendation will be a part of the scheduling interface, but will also be accessible in addition to the scheduling interface.

C. CNII user survey after UMAP 2010

After finishing the UMAP 2010 conference, we conducted a survey in order to gather the opinions of users about the system and to find ways to improve it. This survey consisted of 29 questions, 23 of them with choices from 1-5 (from Strongly Agree to Strongly Disagree). Six of them were open-ended questions.

From a total of 11 users who answered the survey, 100% either agreed or strongly agreed that is a good idea to have a conference scheduling tool and to have it incorporate cross-recommendation in general. When asked if it is useful to see already-scheduled papers in conference timetables– 100% agreed or strongly agreed.

Regarding the recommender, 100% of the surveyed users agree or strongly agree that the presence of content-based recommendations makes the whole system more useful, and 91% agree or strongly agree that it was helpful to highlight recommended papers in the schedule. When asked which of the recommenders was more useful, 27% answered that the content-based recommender was better; 9% thought that the tag-based one was better and 55% of users thought that both of them were useful. Only one person answered that none of the recommenders was useful. One interesting comment given by a user was that the system might need a larger conference to show its potential. Perhaps having only 2 parallel sessions was not enough to show how useful it could be.

These data correlate with the log analysis in confirming that both scheduling and recommendation are very important to the users and that using it for social navigation is useful.

Tagging papers (i.e., adding free user keywords to papers) is another functionality implemented in CNII. Most people agree or strongly agree (72%) that tags are useful in the context of short-lived conference papers, and 63% of users are interested in tagging papers. Additionally, 82% of users think that the ability to find papers through tags is useful. Furthermore, one user who was not interested in tagging papers said, “...the presence of tags did not distract or inconvenience my interaction”.

In CNII, we implemented different groups or communities in order to help people share papers and feedback in more specific subject areas. The idea of “communities” in CNII seemed to be too obscure to users, since 45% of users disagree

or strongly disagree when asked whether the community mechanism was clear. When asked about being interested in contributing papers to the communities, 36% neither agree nor disagree, 9% disagree and 54% agree or strongly agree. Although this result can show that people are hesitant about sharing with a community in CNII, 82% replied positively (agree or strongly agree) when asked if it was useful to see a list of papers collected by the community. This same idea is supported by the 91% of users who think that rating papers under the community (which is called “karma” in CNII) is useful. Under the open-ended question on “how to make community mechanism better,” one user explicitly said that how to/why to contribute to communities was not obvious.

These results also correlate with the log analysis. As we see, both the tagging and community mechanisms were considered quite useful, although not as much as scheduling and recommendation. Yet, users see less reason to contribute to these kinds of feedback, not just use it. Engaging the users to do more tagging and community feedback remains one of the challenges for our project.

Another important feature that we were able to evaluate using the survey was user-awareness. Following the nature of social tagging systems, CNII let the users know who scheduled which paper. This was an important enhancement to the original system and we were interested to see how valuable it was. We found two different results. On one side, 73% of users responded positively when asked if it was useful to see who scheduled which paper, the same 73% replied that it is useful to see co-bookmarking users, and the same 73% agree or strongly agree that it is useful to see the schedules of other users. However, 54% had preferred to keep their schedules private. These results are counterintuitive; however, one of the user’s opinions helped us to understand it: “This could be a situation of not being interested in the paper of someone who would expect –or want– you to be interested in their work. I would rather give false information to a schedule planner to avoid offending someone – such as a PhD starting out, someone who is working in the same field, or someone who is scheduled against a paper of greater interest or relevance.”

V. THE FUTURE: EARLY-START RECOMMENDATIONS (PETER)

While developing the regular recommendation approach for our systems, we were concerned about the known “new user” problem. The nature of this problem is that the profiles of new users, who provided little feedback to the system, are not deemed reliable enough to provide good recommendations. One approach to address this problem is using data from other systems to *fast start* user profiles. This approach is known as mediation [19] or ubiquitous user modeling [20]. We used this approach successfully in our other projects [21] and were interested to explore it in the domain of research talks. In this domain, a natural approach is to build a profile of user research interests, which can be extracted from one of the paper sharing systems such as CiteULike.

To explore this approach, we enhanced our recommendation algorithm in such a way that it can take into

account information about papers bookmarked by the users in CiteULike. We piloted this approach at UMAP 2010, but were not able to evaluate it due to the low number of users who had CiteULike accounts. However, on the level of subjective feedback, 54% agreed or strongly agreed that using CiteULike profiles to offer fast-start recommendations is a good idea, and the rest neither agree nor disagree. The extensive exploration of this idea is still in the future for our project. We plan to explore it in the 2010-2011 academic year in both systems. However, to develop the best version of the enhanced algorithm and to estimate its impact, we performed a predictive evaluation of this approach in the user study. The remaining part of this chapter presents several integrated recommendation approaches that we explored and the results of their evaluation.

A. Improving Recommendations using CiteULike

We explored two ways to improve recommendations in CoMeT: 1) by using information about bookmarked papers from CiteULike in addition to the standard use of information about bookmarked talks from CoMeT; and 2) by using tags to better represent information about talks (and user interest) in addition to standard use of text-only information from talk descriptions. In addition, we combined both approaches – i.e., used both kinds of information (descriptions and tags) from both systems.

While fusing information from two systems is relatively straightforward (from the recommender engine’s point of view, a bookmarked talk and a bookmarked paper is simply a bag of words plus a bag of tags), fusing tags and text in both item representations and profiles is not obvious and can be done in several ways. We start the presentation of our recommendation approaches by introducing several representation models, which explored various ways to fuse keywords and tags. After that, we explain the user profiling and recommendation approaches based on these representations.

B. Document Representation Models

There are many ways to combine various sources of information to build user profiles. We used tags and keywords which appeared in abstracts and titles of CoMeT talks and CiteULike papers. To construct user profiles, we used the following models:

Keywords Only (KO): This is the baseline model described in the previous section. All CiteULike papers can be represented as a $k \times l$ matrix D_c (k is the number of CiteULike papers and l is the number of keywords used in those papers) and all CoMeT Talks are represented in an $e \times m$ matrix D_t (with e as the total number of colloquia and m as the total number of keywords).

To integrate these two sources of data in this model, we obtain a $(k + e) \times (l + m - o)$ matrix D , where o is the number of keywords common to both the CoMeT and CiteULike systems.

Keywords+n*Tags (KnT): In this model, tags are considered to be regular keywords and we treat each document

as a bag of words containing a document’s abstract, title and tags. Each tag appears n times in this bag of words. Each document is represented as a vector in the keywords and tags vector space weighted by the TF.IDF weighting scheme ($d_c = (w_{1,c}, w_{2,c}, \dots, w_{l+j-p,c})$ where $w_{i,c}$ shows the weight of i^{th} keyword in document c , l is the total number of keywords, j is the total number of tags, and p is the number of common terms between tags and keywords).

In this case, we can create a merged documents matrix just as in the previous model.

Keywords Concatenated by Tags (KCT): In this model, we consider tags to be a separate source of information, and we treat them separately. We obtain a bag of keywords and a bag of tags specific to each document. Using TF.IDF weighting scheme, a tags vector and a keywords vector is built for each document. Next, each document is represented by concatenating keywords and tags vectors as one vector in the keywords and tags vector space ($d_c = (w_{1,c}, w_{2,c}, \dots, w_{l,c}, t_{1,c}, t_{2,c}, \dots, t_{j,c})$ where $w_{i,c}$ shows the weight of i^{th} keyword in document c , l is the total number of keywords, $t_{i,c}$ shows the weight of i^{th} tag in document c , and j is the total number of tags).

In this case, we will have an $e \times (l + j)$ CoMeT talk’s matrix (D_t) where e is the number of CoMeT talks, l is the total number of keywords in CoMeT, and j is the total number of tags in it. We will also have a $k \times (m + i)$ CiteULike papers matrix (D_c) where k is the number of CiteULike talks, m is the total number of keywords in CiteULike, and i is the total number of tags in it. After merging these two matrixes, we will have an $(e + k) \times (m + l + j - o - p)$ matrix D , showing all documents in keywords and tags vector space in which o and p are respectively the number of common keywords and tags between two systems.

To study the impact of various sources of information on recommendation systems, we utilized each of the models described above with only CoMeT’s data, and another time with CiteULike’s data sources.

C. The evaluation of the new approach

In this section, we provide experimental results of adding CUL profiles to the CoMeT recommendation system. In this study, the top 10 recommendations of each model are given to CoMeT users possessing CUL accounts and these users were asked to answer three questions measuring the relevance of the talk to research interests, overall interest, and novelty:

- Is this talk related to your interest? (yes/no question)
- How interesting is this talk to you? (in 5-point scale)
- If the talk is related to your interests, how novel is this talk to you? (in 5-step scale)

We compared models presented in the previous section (with $n = 1, 2, 5$ for KnT model); one time with only CoMeT user profiles and the other time with both CoMeT and CUL user profiles. The $nDCG$ measure is calculated for each model to compare how each model can recommend more interesting talks to users (Table V). The precision measure is calculated (as shown in Table IV) to compare the ability of each model to

provide relevant talks to users (based on the results of yes/no questions). To reduce the volume of reported data, in all tables, we only show results for the best n value for KnT model, which was $n=1$.

As we can see in Table IV and Fig. 12, adding tags using the fusion approach (KnT), results in better cumulative precision for the top 10 recommendations. The exceptions are the very top positions in the recommendation list, where the KO model works better in both cases. This interesting effect is caused by the different behaviors of these two models. For the KO model, the precision is high at the top positions, but then drops rapidly. The precision of KnT model is more stable and surpasses the KO model at position 5.

TABLE IV

PRECISION RESULTS FOR DIFFERENT MODELS WITH DIFFERENT NUMBERS OF RECOMMENDATIONS

Precision		1	2	3	4	5	6	7	8	9	10
Only CoMeT Data	KO	.83	.67	.72	.63	.6	.56	.57	.5	.51	.5
	KnT $n=1$.5	.5	.58	.59	.57	.58	.57	.58	.6	.6
	KCT	.5	.33	.39	.46	.47	.53	.52	.5	.5	.5
CoMeT + CiteULike Data	KO	.83	.83	.67	.75	.73	.69	.64	.63	.56	.56
	KnT $n=1$.63	.69	.71	.72	.73	.73	.71	.7	.68	.68
	KCT	.38	.44	.42	.47	.48	.52	.5	.49	.53	.53




Fig. 12. Precision Results for different models with different numbers of recommendations.

Incorporating CoMeT data into both the KnT and KO models increases the precision dramatically. This proves that tags within KnT and CoMeT can help with the precision. Moreover, these two effects seem to be stackable: the KnT model which includes both CoMeT and CiteULike (CUL) data has the best cumulative precision.

TABLE V

NDCG RESULTS FOR DIFFERENT MODELS WITH DIFFERENT NUMBERS OF RECOMMENDATIONS

nDCG		1	2	3	4	5	6	7	8	9	10
Only CoMeT Data	KO	.9	.88	.89	.93	.92	.94	.95	.95	.95	.96
	KnT $n=1$.9	.85	.82	.83	.87	.88	.89	.9	.91	.93
	KCT	.84	.88	.89	.9	.9	.91	.92	.92	.94	.95
CoMeT + CiteULike Data	KO	.84	.91	.9	.92	.93	.94	.95	.96	.96	.96
	KnT $n=1$.9	.9	.89	.88	.9	.92	.92	.94	.94	.95
	KCT	.77	.85	.84	.81	.83	.84	.86	.88	.91	.92




Fig. 13. Interest results (nDCG) for different models with different numbers of recommendations.

At the same time, adding tags via the KCT model degrades the system’s precision. This might be because of high dimensionality of our vector space model when we concatenate the keywords vector with the tags vector. In this case, the distance between documents and user profile increases and decreases the variance between similarities of user profile to different talks. We can see that different sources of information may result in less precise results if we don’t integrate them in an appropriate way.

The results for “interest” in Table V and Fig. 13 are a bit different. Generally, we can see the positive effect of using CiteULike data. In both KCT and KnT models, using CiteULike data in addition to CoMeT data boosts user cumulative interest in returned talks, although this difference is very small for the KO model. At the same time, the best results are produced by the tag-less KO model both with and without CiteULike data. These results are very close to each other and better than other models in general.

To explain these results, we should stress that we deliberately separated “relevance”, which was understood as the fit with a user’s research work and “interest”, which was considered to be part of the overall attractiveness of an item. Through the CoMeT experience, we observed that many users seem to be interested in some talks on general topics (like art and politics), which had little in common with their research interests. A separation of relevance and interest allowed for cases where a talk is rated as interesting, yet non-relevant. The analysis of user rating data confirmed that we were correct – there were a number of talks like that for almost all users. We think that the decrease in the system’s ability to recommend *interesting* talks with the addition of tags can be caused by the increased focus of relevance encapsulated in tags. It naturally decreased the system’s ability to recommend interesting, but not relevant talks. This is a natural outcome of user tagging behavior, which was focused mostly on their research interests. As a result, a simpler KO model is able to better grasp the user’s overall interests.

Users usually employ CiteULike to add, review and rate papers related to their research field. This requires a user to spend a considerable amount of time on a paper and users prefer to review only those papers related to their field of research. On the other hand, CoMeT contains information about talks happening within a specific time given on a particular date. It is more plausible for a user to bookmark a

more interesting, less relevant talk knowing that he/she might miss this amount of information given in a limited time. As a result, CoMeT user profiles include wider areas of user interests as compared to CiteULike user profiles, which usually contain more relevant documents. Comparing the results, we can see that there are several trends. First, the addition of tags using the KnT fusion model helps to improve relevance of results. In contrast, when a user's general interests are considered, the tag-less KO model produced slightly better results. As for KCT concatenation model, it seems to have real problems: it decreases system performance for both kinds of measures producing poor performances in most cases.

The effect of adding CiteULike is more consistent. It typically produces better results for all measures, although its effect on measuring interest is negligible. Interesting enough is that the effects of adding tags and adding data appeared to be stackable. i.e., an approach, which uses both tags and CiteULike "stacks" the separate effects of the component approaches resulting in best approaches for relevance and novelty measures.

VI. CONCLUSION

This paper presented the results of our work on two different social systems for sharing research talks – Conference Navigator and CoMeT. Both systems focused on the problem of building successful social systems for smaller communities. The main problem that we attempted to handle is the low volume of contributions, which endangers social systems which have smaller number of users. To make the systems more useful and to encourage the users to contribute, we used several ideas such as integrating social navigation and social tagging features as well as using the "personal need" approach to increase the volume of contributions. Our early results indicated that our ideas were successful in reaching a higher-than-normal rate of contributions, but that such a rate was still not sufficient in cases where a personal schedule – our main "personal need" idea – is not critical to the users. In our current work, we are attempting to combine personal scheduling with personal recommendation, another "personal need" approach. The new results indicate that this approach could support further growth in scheduling types of contributions. Among the challenges left for future research is to increase the precisions of recommendations and to provide incentives for user tagging and community contribution activities.

REFERENCES

- Ling, K., Beenen, G., Ludford, P., Wang, X., Chang, K., Li, X., Cosley, D., Frankowski, D., Terveen, L., Rashid, A.M., Resnick, P., and Kraut, R.: 'Using social psychology to motivate contributions to online communities', *Journal of Computer-Mediated Communication*, 2005, 10, (4), pp. article 10
- Nielsen, J.: 'Participation inequality: encouraging more users to contribute', in Editor (Ed.)^(Eds.): 'Book Participation inequality: encouraging more users to contribute' (2006, edn.), pp.
- Lui, S., Lang, K., and Kwok, S.: 'Participation Incentive Mechanisms in Peer-to-Peer Subscription Systems', in Editor (Ed.)^(Eds.): 'Book Participation Incentive Mechanisms in Peer-to-Peer Subscription Systems' (IEEE Computer Society, 2002, edn.), pp. 302-309
- Cheng, R., and Vassileva, J.: 'Design and evaluation of an adaptive incentive mechanism for sustained educational online communities', *User Modelling and User-Adapted Interaction*, 2006, 16, (2/3), pp. 321-348
- Cheng, R., and Vassileva, J.: 'User- and community-adaptive rewards mechanism for sustainable online community', in Editor (Ed.)^(Eds.): 'Book User- and community-adaptive rewards mechanism for sustainable online community' (Springer Verlag, 2005, edn.), pp. 332-336
- Hummel, H.G.K., Burgos, D., Tattersall, C., Brouns, F., Kurvers, H., and Koper, R.: 'Encouraging contributions in learning networks using incentive mechanisms', *J. Comput. Assist. Learn.*, 2005, pp. 355-365
- Shapira, B., Kantor, P., and Melamed, B.: 'The effect of extrinsic motivation on user behavior in a collaborative information finding system', *J. Amer. Soc. Inform. Science & Techn.*, 2001, 52, (11), pp. 879-887
- Bretzke, H., and Vassileva, J.: 'Motivating cooperation on peer to peer networks', in Editor (Ed.)^(Eds.): 'Book Motivating cooperation on peer to peer networks' (Springer Verlag, 2003, edn.), pp. 218-227
- Bhattacharjee, R., Goel, A., and Kollias, K.: 'An Incentive-Based Architecture for Social Recommendations', in Editor (Ed.)^(Eds.): 'Book An Incentive-Based Architecture for Social Recommendations' (ACM, 2009, edn.), pp. 229-232
- Farzan, R., Dimicco, J., Millen, D., Dugan, C., Geyer, W., and Brownholtz, E.: 'Results from deploying a participation incentive mechanism within the enterprise', in Editor (Ed.)^(Eds.): 'Book Results from deploying a participation incentive mechanism within the enterprise' (ACM, 2008, edn.), pp. 563-572
- Speer, R., Krishnamurthy, J., Havasi, C., Smith, D., Lieberman, H., and Arnold, K.: 'An Interface for Targeted Collection of Common Sense Knowledge Using a Mixture Model', in Editor (Ed.)^(Eds.): 'Book An Interface for Targeted Collection of Common Sense Knowledge Using a Mixture Model' (2009, edn.), pp. 137-146
- von Ahn, L., and Dabbish, L.: 'Designing games with a purpose', *Commun. ACM*, 2008, 51, (8), pp. 58-67
- Farzan, R., and Brusilovsky, P.: 'Social navigation support in a course recommendation system', in Editor (Ed.)^(Eds.): 'Book Social navigation support in a course recommendation system' (Springer Verlag, 2006, edn.), pp. 91-100
- Farzan, R., and Brusilovsky, P.: 'Social navigation support through annotation-based group modeling', in Editor (Ed.)^(Eds.): 'Book Social navigation support through annotation-based group modeling' (Springer Verlag, 2005, Edinburgh, UK edn.), pp. 463-472
- Farzan, R., and Brusilovsky, P.: 'Where did the Researchers Go? Supporting Social Navigation at a Large Academic Conference', in Editor (Ed.)^(Eds.): 'Book Where did the Researchers Go? Supporting Social Navigation at a Large Academic Conference' (2008, edn.), pp. 203-211
- Claypool, M., Le, P., Wased, M., and Brown, D.: 'Implicit interest indicators', in Editor (Ed.)^(Eds.): 'Book Implicit interest indicators' (ACM Press, 2002, edn.), pp. 33-40
- Wongchokprasitti, C., Brusilovsky, P., and Parra, D.: 'Conference Navigator 2.0: Community-Based Recommendation for Academic Conferences', in Editor (Ed.)^(Eds.): 'Book Conference Navigator 2.0: Community-Based Recommendation for Academic Conferences' (2010, edn.), pp.
- Freyne, J., Farzan, R., Brusilovsky, P., Smyth, B., and Coyle, M.: 'Collecting community wisdom: Integrating social search & social navigation', in Editor (Ed.)^(Eds.): 'Book Collecting community wisdom: Integrating social search & social navigation' (ACM Press, 2007, edn.), pp. 52-61
- Berkovsky, S., Kuflik, T., and Ricci, F.: 'Meditation of user models for enhanced personalization in recommender systems', *User Modeling and User-Adapted Interaction*, 2008, 18, (3), pp. 245-286
- Carmagnola, F., and Dimitrova, V.: 'An Evidence-Based Approach to Handle Semantic Heterogeneity in Interoperable Distributed User Models', in Editor (Ed.)^(Eds.): 'Book An Evidence-Based Approach to Handle Semantic Heterogeneity in Interoperable Distributed User Models' (Springer Verlag, 2008, edn.), pp. 73-83
- Sosnovsky, S., Brusilovsky, P., Yudelso, M., Mitrovic, A., Mathews, M., and Kumar, A.: 'Semantic Integration of Adaptive Educational Systems', in Kuflik, T., Berkovsky, S., Carmagnola, F., Heckmann, D., and Krü ger, A. (Eds.): 'Advances in Ubiquitous User Modelling' (2009), pp. 134-158