Improving Scientific Conferences by Enhancing Conference Management Systems with Information Mining Capabilities

Andreas Pesenhofer
iSpaces Group,
eCommerce Competence Center – EC3
Donau-City-Straße 1, A–1220 Wien, Austria
Email: andreas.pesenhofer@ec3.at

Rudolf Mayer and Andreas Rauber
Department of Software Technology and Interactive Systems,
Vienna University of Technology
Favoritenstraße 9–11/188, A–1040 Wien, Austria
Email: {mayer, rauber}@ifs.tuwien.ac.at

Abstract

In this paper we identify tasks in the field of conference management where methods from the domain of information retrieval, information management and information organization can assist the organizer, the program committee member and the participants. In particular, we focus on tasks, where (1) the quality of the conference can be increased by assisting in the creation of an improved review process by better matching the reviewers expertise with the paper topics, (2) the conference participants profit by allowing them better access to the wealth of information accumulated throughout a conference series and at the same time (3) reducing the workload of conference organizer by partially automating tedious tasks, such as review assignment and the creation of review plans. We report on case studies from a medium-sized (around 400 participants) as well as a large (more than 700 participants) conference in computer science as well as in medical domains.

1. Introduction

One of the latest released reports from the International Congress & Convention Association (ICCA)[1] deals with the statistics of organized meetings for the year 2005. These rankings cover meetings organised by international associations which take place on a regular basis, have more than 50 participants and rotate between a minimum of three countries. For the year 2005 the ICCA Data researchers have identified 5,315 events, a rise of 511 over 2004. These statistics correspond with the list of 455 forthcoming conferences (mostly announced via DBWorld) for the timeframe May 3rd 2006 to December 19th 2007[2].

The organization of a scientific conference is a challenging endeavor where a small error can have tremendous influence on the event. The IEEE, for example, provides a conferences organization manual[3]. For the technical (scientific) part of the conference the use of web-based management system (such as [2,7,8,10–12]) is indispensable in handling the huge amount of submissions. These systems fulfill the basic requirements and drastically ease organization. Yet, there are still many tasks where methods from the domain of information management and information visualization can assist to further improve the quality as well as to reduce the workload.

This paper describes the tasks in a conference management system where the use of information mining capabilities provides advanced methods to assist the organizer, the program committee member and the participants. The core goals of the work presented in this paper are (1) to increase the quality of the conference by assisting in the creation of an improved review process by better matching the reviewers expertise with the paper topics; (2) to increase the benefits to conference participants by allowing them better access and utilization of the wealth of information accumulated throughout a conference series, while at the same time

[1] http://www.iccaworld.com

[2] http://dbms.uni-muenster.de/menu.php3?item=confs

[3] http://www.ieee.org/web/conferences/mom/

(3) easing the workload of conference organizers by partially automating tedious tasks, such as review assignment and the creation of review plans.
The remainder of this paper is structured as follows. Section 2 gives an overview of related work. Section 3 describes the basic functionalities of a conference management system. Three core tasks for further automatization will be tackled in Section 4, Section 5 and Section 6. Finally, we summarize in Section 7.

2. Related Work

Conference management systems are web-based systems that cover some tasks so that the organization of scientific conferences can be carried out a little bit easier. Such tasks are for example the collection of submissions, the handling of assigned papers that the Program Committee (PC) members have to review, the download of papers, the handling of reviewers preferences and bidding, review progress tracking, web-based PC meeting, notification of acceptance/rejection and sending e-mails for notifications to authors or PCs. Once a bidding process has been performed, the assignment is handled as an optimization problem, trying to allocate papers according to reviewer preferences while striving for equal load distribution.
Dumais and Nielsen [3] used data given by 15 reviewers that consisted not only of the submitted abstracts and/or interests, but also provided complete relevance assessments for the 117 submitted papers. Information retrieval principles and latent semantic indexing were used to generate the automatic assignments for each reviewer. This method achieved improvement of 48% compared to the random assignment.
Yarowsky and Florian [13] focused on the classification of every paper to exactly one of six conference committees. They used 92 papers which were submitted to the ACL conference in electronic form and additionally requested committee members to provide representative papers so that a reviewer profile could be created. First a centroid for each reviewer and then a centroid for each committee as the sum of its reviewer centroids was computed. For each paper the cosine similarity was computed and compared with the committee centroids where the highest rank was the selection criteria. They concluded that the automatic methods could be as effective as human judges, especially in case where the judges may be less experienced.
In [9] the assignment of papers is done based on previous collected user ratings. The paper describes a simple method which provides an approximate solution to the problem without requiring each user to rate each item. The method relies on an interactive process where in each step (or ballot) the users have to rate a sample of items. Collaborative filtering is then performed to predict the missing ratings as well as their level of confidence. Performing a new ballot may improve the accuracy of the prediction. This algorithm tends to lead to a suboptimal solution if only a sub group of reviewers rates the ballot and if only one ballot round is performed.
In [10] the assignment is made based on the bids for special papers and on the the reviewers' expertise on the conference topics and the willingness to review papers on these topics. The reviewers may bid in several stages and the bids are accumulated. Graph theory is applied to carry out the assignment.
The latest work in this domain was carried out by Aleman-Meza et al. [1] where they describe a semantic web application that detects conflict of interest relationships among potential reviewers and authors of scientific papers. The degree of conflict of interest between the reviewers and authors is calculated based on a populated ontology. As input they integrated entities from two social networks, namely 'knows' from a FOAF (Friend-of-a-Friend) social network and 'co-author' from the underlying co-authorship network of the DBLP bibliography. This allows them to detect more potential conflict of interests than the simplified method that is implemented in [8].

3. Conference Management

3.1. User Roles and Tasks

In a conference management system users with different roles have to have access to specific tasks in a predefined time slot. An analysis of these roles and tasks is given in [4] and [8]. We have to distinguish between organizers, PC members or reviewers, authors, participants and persons visiting the web page. The program committee (PC) chair is in charge of the coordination and monitoring of the necessary tasks.
Such tasks include setup/customization, paper submission, conflict of interest detection, reviewer assignment, reviewing, paper selection, session creation, poster setup plans and conference participant support. In this paper we will concentrate on conflict of interest detection together with the automatic assignment of submissions to reviewer as well as the creation of poster setup plans and the conference participant support.

3.2. Tasks for further automatization

In this section we will focus on tasks where further automatization eases the work of PC members and the PC chairs. We furthermore try to identify means to assist conference participants both at as well as (particularly) after the conference in order to make the most of the wealth of infor-

mation presented during the meeting and accumulated over the years in a conference series.

3.2.1 Task: Reviewer Assignment

The submission to reviewer assignment is done either automatically or manually by the PC chair, with an automatic assignment usually being followed by a manual adjustment. For the assignment the following constraints are taken into consideration:

- The submission topics should match with reviewer interests.
- A reviewer's bid for specific papers has to be taken into consideration.
- Reviewers should not get their own paper to review. A potential conflict of interest between the PC members and submissions has to be calculated.
- Each PC member should get the same amount of papers to review, so that they have the same work load.

All these tasks rely on the input of the PC members. This can cause trouble if some of the PC members are reluctant or too busy. It is not possible for the PC chair (administrator of the system) to make decisions for them, being limited to sending reminder mails and asking for their cooperation. The algorithm can not work properly and produces suboptimal solutions that have to be corrected manually by the PC chair.
In Section 4 we will focus on an automatic assignment of the submitted papers to the PC members based on their previous publications as a baseline for the manual bidding process. It overcomes these problems by using publicly available publications of the authors to create the PCs' profiles.

3.2.2 Task: Poster Alignment

During most of the conferences posters are presented in a special room or in the lounges of the conference venue. Usually there exists a pre-setup provided by the organizer where authors have to fix their posters. In this case the organizers have to figure out which posters fit best together when grouped by topic. The PC chair has to align the poster manual. Mnemonic SOMs as described in Section 5 can be used for this alignment.

3.2.3 Task: Participant support

The conference program should be kept up to date in the web and the proceedings should be searchable either publicity or limited to registered conference participants via dedicated logins. Participants may be interested if they have missed interesting sessions. Mnemonic SOMs and SOMs in combination with the participant's interests give the participants new insight in the huge amount of information presented during the conference as well as helping them to prepare their schedule before attending large events. We will address that in more detail in Section 6.

There are numerous tasks that offer considerable potential for automatization, such as in the production of consistent printing and on-line material (e.g. web, program broshure, proceedings, notice boards), accounting, etc. which are not dealt with in this paper.

3.3. Case Studies

We report on case studies from two conferences, the *9th European Conference on Research and Advanced Technology for Digital Libraries* (ECDL 2005)[4] and the *European Congress of Radiology 2004* (ECR 2004)[5]. The ECDL is the major European conference on digital libraries and associated technical, practical and social issues in this field. It can be classified as a mid sized conference with around 100 to 200 submissions, around 80 to 90 program committee members and around 350 to 450 participants. The ECR is a large sized conference species with more than 2,000 scientific paper submissions taking place every year in Vienna. It is the largest radiological meeting in Europe attracting more than 15,000 participants from over 90 countries. WEBGES[6], who is the soft- and hardware provider of the ECR, provided us with the relevant data.

The data has to be transformed into a representation so that it is understandable by the algorithms. Therefore, we indexed the data based on the the well known bag-of-words approach with Lucene using a tfidf weighting scheme, which is based on the term frequency (tf) in the given document and the inverse document frequency (idf) of the term in the whole collection. Pre-processing steps in form of removing all numbers, punctuation marks and special characters were applied. The tfidf values were normalized to unit vector length.

3.3.1 ECDL corpora

For the ECDL we have to distinguish between three corpera:
ECDL A: Is made out of 723 automatically retrieved publications from PC member's home pages and 125 submissions. Term reduction based on document frequency and term length was applied, resulting in a vector with 8,767 unique terms.
ECDL B: Consists of the accepted poster submissions, 30 different posters in the English language. After applying an

[4] http://www.ecdl2005.org
[5] http://www.ecr.org
[6] http://www.webges.com

English stop word list and other term reductions based on the document frequency, we obtained a vector of 569 different terms.
ECDL C: Is composed at the accepted paper and poster submissions, totaling to 71 documents. Applying the same mechanisms as for the ECDL B corpus, we obtained a vector of 5,654 different terms.
In all three cases no stemming was applied.

3.3.2 ECR corpus

This corpus consists of the abstracts of the ECR from the year 2004. All together there are 943 English documents which were presented during the scientific sessions of the congress and which each belong to one of the 15 different topics (c.f. Table 3). Every abstract is assigned to exactly one topic. Additional to the settings that were described in the beginning of the section we also applied an english stop word list and only kept those terms that had a df between two and 300. In the end, the corpus consisted of 3,842 unique terms.

Additionally we received the radio frequency identification (RFID) logs that were collected during the conference. At the registration every participant received a badge with a unique RFID tag. The entrances to halls of the conference location were guarded with RFID gates, so that the organizer could track access to a session. These are used in the medical domain for the monitoring and issuing of continuous education certificates. They serve to build an anonymized participant profile for our experiments.

4. Profile based Reviewer Assignment

A good paper to reviewer assignment is based on the cooperation of the PC member (reviewer). They have to choose from a list of relevant topics which they are interested in and furthermore they have to bid for special papers by skimming through the abstracts. Most of the PC members neither bid nor choose their interests so that the algorithms fail in computing a proper assignment. This is particularly due to the fact that a bidding process for 200 or more papers is a notoriously time consuming task. Our solution overcomes this problem, because the interest of the reviewer is defined based on previous publication that are available on the internet.

4.1. Profile generation

We used the forename and the surname of the PC members to formulate the search query. We sent the query to two search engines which provide scientific papers, namely CiteSeer.IST and GoogleScholar. From the returned search result pages the URLs linking to the publications were extracted. Using the 87 PC members from the ECDL 2005 conference resulted in 4,369 retrieved URLs. In the next step we downloaded these documents discarding all non PDF documents. As result we obtained the 723 potential publications. Note that for ten PC members no publications have been automatically retrieved.

4.2. COI detection

The potential conflict of interest detection (COI) was performed based on (1) the occurrence of the last name of a program committee member in the authors line of a submission and (2) the existence of parts from the PC members email domain in the submissions author field (e.g. if the PC members email has the domain tuwien.ac.at and the submission comes from the same domain then we will register a potential conflict of interest). Using these two methods allowed us to identify 46 potential conflict of interest for the PC members for the ECDL 2005 data set.

We compared our results with the COI that the PC members registered during the biding phase of the ECDL conference. Here in only 24 cases a COI was registered. A detailed comparison of the two lists reveals the following:

1. More than the half (57.69%) of the reviewers did not bother to register a COI. This group of people was additionally identified by our system.

2. Potential COI was detected by the system but not registered by the reviewer, who in principal did register 50% of the COI. As reasons we identified that the COI was not considered in spite of being from the same lab, because of a lack of close cooperation and that the paper was overlooked due to the large list of papers. A solution would be to have a system that detects a potential COI and presents it to the reviewer to confirm it.

3. In 7 cases the COI was registered by the reviewer, but not detected by our current system. In these cases coauthorship analysis would have to be included (e.g. DBLP) and for areas that are not covered by a specific digital library of papers a web-based search has to be performed.

4.3. Reviewer Assignment

Before we can calculate the assignment, we have to find out which submissions match with the interests of which PC member. Therefore, we computed the Euclidian distance between every submission and publication based on the full-text indexed feature vector. A distance of 0 means that the two compared documents are identical and the higher the

Table 1. Distribution of the review workload

paper/reviewer	preferences & bid-based	profile-based
4	40	12
5	9	27
6	2	25
7	5	17
8	31	6
PC member (sum)	87	87

value is the more different they are. A PC member has normally more than one publication in his profile, so we kept only the smallest distance from all his documents to one submission.

As baseline for our evaluation we used the automatic assignment that was calculated on the ECDL 2005 PC member preferences and their bids. To make our system comparable with the baseline we set up an identical system without the bids and the paper topic interest of the PC members. The aggregated distances were sorted starting with the smallest and ending with the largest. The first ten received a rate level of 4 which correspond to a bid of 'eager' to review, the next ten were rated with 'interesting' (3) and the remaining received the level 1 ('better not'). For the ten cases where no publications could be found automatically, and therefore no distances to the submission existed we used 2 ('indifferent') as default rating. If a COI in the relation was detected a rate level of 0 ('conflict of interest') was inserted into the data base. These pre-calculated values serve as a basis for the bidding process that may be optimized by the user.

Table 1 summarizes the workload distribution of the PC members using the assignment model based on preferences and bids compared to the results that where obtained with the profile-based assignment. In both cases we have 500 reviews that have to be assigned to the 87 PC members, the optimal amount of assigned papers per PC member would have been 5,75. In the first case, the preferences & bid-based model, 40 reviewers get four papers to review and 31 reviewers get the maximum amount of papers (8) assigned. Only 16 reviewers get 5 to 7 papers assigned. In our system, the profile-based one, only 6 reviewers have a workload of 8 papers and 12 PC member have only four papers to review. Most of the PC members (27) got 5 papers, followed by 25 that got 6 and 17 that got 7 papers assigned. In this case many more PC members are allocated around the mean of 5,75 resulting in a more equal distribution than in the first case.

5. Poster setup plans

When the setup of the poster locations is defined by the conference organizers it can be done in one of several different ways, for example, the setup may be organized completely randomly or sorted alphabetically by author names or submission titles. It may be desirable, though, to organize the submissions by their content - that way, conference participants can easily find the areas with posters about topics they are interested in. Organization by content may be done using manually assigned category labels coming either from the authors themselves during submission, or from the PC. However, such a categorization may in many cases not be available at all, available only for some parts of the submissions or of poor or varying quality. Then, as an alternative, unsupervised clustering algorithms based only on the submission contents may be utilized to determine a poster setup.
Independent of the exact setup, the conference participants should also be provided with a map of the venue, indicating poster locations and topic areas, in order to assist them in locating the posters they are interested in.

Unsupervised clustering and generating a map of the poster setup can be achieved using for example the Self-Organizing Map (SOM) [5]. The SOM is a neural-network model that provides a mapping from a high-dimensional input space to a lower dimensional output space. In this mapping, the SOM preserves the topology of the input space, i.e. input patterns that are located close to each other in the input space will also be located closely in the output space, while dissimilar patterns will be mapped on to opposite map regions. In many applications, this output space is made of a two-dimensional, rectangular map. This representation allows for an easier interpretation of the complex structure of the input patterns by the user.

Another advantage of using the SOM is that it generates a clustering that preserves transitions between clusters - documents that would belong to two different clusters will be mapped on the border in between those clusters.

In our application, the input space will be formed by a vector-space representation of the poster submissions, as described in Section 3.3, while the output space will be the map of the poster session area.

As in many cases, the area for the poster session may not be of rectangular shape, we use a modifcation to the original SOM algorithm, the *Mnemonic SOM*, as presented in [6]. In the Mnemonic SOM, the output space is two-dimensional, but can take any arbitrary shape. They can be easily generated from a black and white image representing the desired shape, for example for the poster presentation area.

We have applied this method for arranging the poster setup during the *9th European Conference on Research and*

Table 2. Class legend for the ECDL 2005 data

Concepts of Digital Libraries, Documents and Metadata		Digital Preservation, Web Archiving	
System Architectures, Open Archives, Integration		Digital Library Applications & Case Studies	
Information Retrieval & Organization, Search & Usage		Multimedia, Audio, Video	
User Studies, Evaluation, Personalization, UI			

Advanced Technology for Digital Libraries (ECDL 2005). Table 2 gives an overview of the topics of the submitted posters to this conference. The category assignment was given by the authors on submission.

Figure 1 shows the generated mapping, where the output space was made of a grid with the size of 35x15, with 182 units within the map shape. It is based on the layout of the conference poster area. Black lines show the setup of the poster boards, and numbers indicate the ID assigned to each poster on submission.

We can observe that thematically similar posters get arranged close to each other, for example in the top-left we can find posters dealing with the 'Open Archives Initiative Protocol'. The poster arrangement does not necessarily follow the manual categorization, but arranges them by content.

The given data set contains a lot of different, sometimes rather small clusters. This is due to the small size of the data set (30 accepted posters), and the very heterogenous topics they discuss. However, the quality of the generated mapping is good.

Using the method described above can help the conference organizer both in saving time on the poster setup and in achieving a better thematically grouped setup.

6. Participant support

The method of the SOM, described in Section 5, can also be well utilizied for supporting the participant during and after the conference.

One application is to provide an advanced interface to the proceedings of the conference, in addition to traditional key-word based searching or manually created indices. We again generate representations of all the presentations at the conference via a vector-space represantion of the abstracts, and map the documents on a SOM. Figure 2 gives an example from the ECDL 2005 conference in Vienna (cf. Section 5), where we use a map in the shape of Austria as a mnemonic hint for the participants. The submissions, including both papers and posters, have been grouped automatically according to their topic by the SOM algorithm. The colored pie-charts visualize the distribution of the manually assigned categories of the docuements. The labels on the map (e.g. Query, Search, Machine Learning) have been added manually after inspecting the content of the documents grouped together in this region.

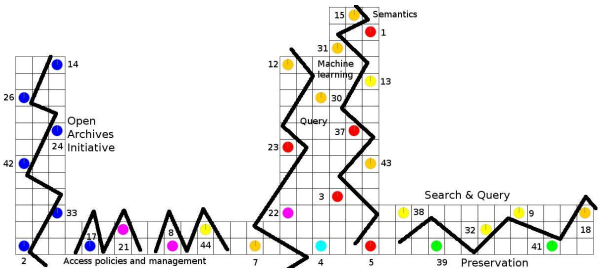

Figure 1. Poster allignment for the ECDL 2005 conference.

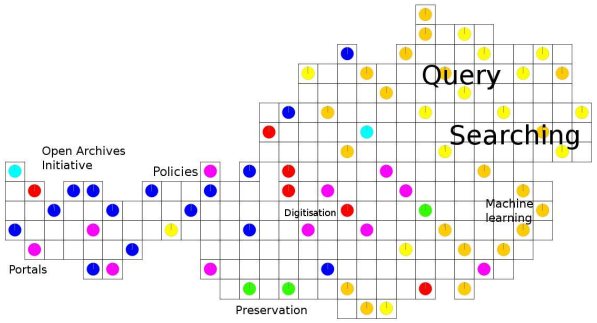

Figure 2. Map of the submission to ECDL 2005.

The scientific abstracts of the ECR 2004 were also preprocessed as described in Section 3.3.2 and mapped onto a SOM, this time following the shape of the logo of the Austria Center Vienna (ACV), the location where the conference takes place every year. The shape of the logo also represents the basic form of the ACV building.

Figure 3 illustrates how this content based mapping on the shape of the ACV is done by the SOM algorithm. In Table 3 the category names and colors of the ECR 2004 are provided, so that the evaluation of the map with the colored pie-charts, can be done easily. On the far left corner papers dealing with 'Vascular' (magenta; mark 1) are arranged together. The papers dealing with 'Computer Applications' (orange; mark 2) have their cluster on the right hand side. Papers dealing with 'Interventional Radiology' (grey) are

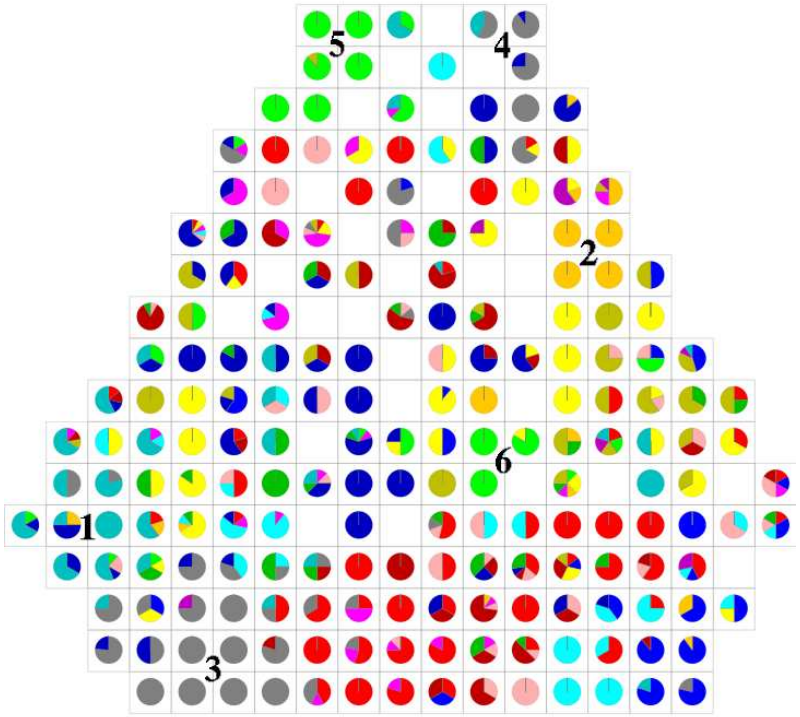

Figure 3. Scientific submissions to the ECR 2004 mapped on the ACV logo

Table 3. Class legend for the ECR 2004 data.

Abdominal and Gastrointestinal		Interventional Radiology	
Breast		Musculoskeletal	
Cardiac		Neuro	
Chest		Pediatric	
Computer Applications		Physics in Radiology	
Contrast Media		Radiographers	
Genitourinary		Vascular	
Head and Neck			

split up into two clusters, where the first one (mark 3) deals with embolization and the second one (mark 4) deals with different kinds of stents. In the neighboring cluster (mark 5) the papers also deal with stents, in particular with coronary artery stents, chest pain and thrombus detection belonging to the class ‘Cardiac’ (light green). In the second ‘Cardiac’ region (mark 6) the documents deal with ventricles, myocardial infarctions and myocardial scars.

In Figure 4 we used the attendance information of a participant (RFID logs) to create personalized fingerprints. We identified the locations of these abstracts that were presented in sessions that the participant attended and created a hit histogram. The more focused a participant is, the more concentrated the histogram appears on the map. Participants can immediately see where their interests are located on the map and by looking at the regions surrounding their fingerprints they may find relevant information to them.

In Figure 4(a) shows a participant interested in ‘Interventional Radiology’ (lower left side) and ‘Neuro’ (top left side).

For the second participant (Figure 4(b)) six regions are highlighted. The two sessions entitled ‘Myocardial viability and wall-motion’ and ‘Evaluation of cardiac function’ both are part of the ‘Cardiac’ topic. They are located next to each other, forming a larger cluster on the right. Two documents from the last session are mapped to the top left of the map, where a second ‘Cardiac’ cluster can be identified. The session dealing with ‘Molecular Imaging’ papers can be found on half way down to the ‘Neuro’ region. Going down and a little bit to the right we come to the ‘Musculoskeletal’ session and going diagonal to the left we end up in the furthermost left spot, described as the ’Interventional Radiology’ section. This participant attended sessions with five different topics, which can be seen by the fingerprint.

In case of the ECR the personalized fingerprint can be added to the profile of the user. The accepted scientific papers of the upcoming conference can be trained as a mnemonic SOM in the shape of the ACV. Using the stored fingerprint

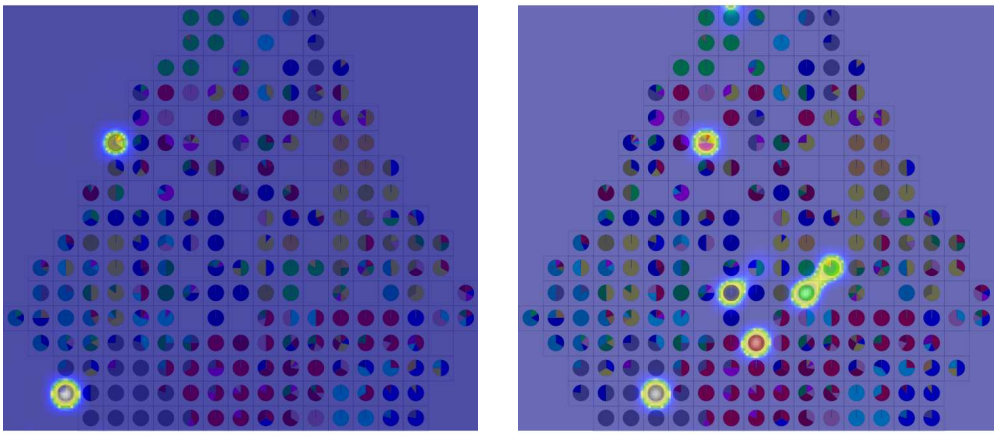

(a) participant A (b) participant B

Figure 4. Fingerprints

allows the participants to mark their interests on the actual conference map and helps them to decide which of the sessions to visit.

7. Conclusion

We presented information mining methods that enhance scientific conference management systems. We showed that organizers, reviewers and participants of mid sized and large sized conferences benefit from our proposed methods which (1) ease the task of paper to reviewer assignment for the organizer, (2) result in more equal paper distribution, (3) help in creating poster setup plans and (4) provide the registered participants better access to the scientific papers and also help them to decide which sessions they should visit at the next conference.

References

[1] B. Aleman-Meza, M. Nagarajan, C. Ramakrishnan, L. Ding, P. Kolari, A. P. Sheth, I. B. Arpinar, A. Joshi, and T. Finin. Semantic Analytics on Social Networks: Experiences in Addressing the Problem of Conflict of Interest Detection. In *Proc. of the 15th International World Wide Web Conference*, 2006.

[2] ConfMan (Conference Manager). [Online]. Available: `http://www.ifi.uio.no/confman/ABOUT-ConfMan/`

[3] S. T. Dumais and J. Nielsen. Automating the Assignment of Submitted Manuscripts to Reviewers. In *Research and Development in Information Retrieval*, pages 233–244, 1992.

[4] P. Halvorsen, K. Lund, V. Goebel, T. Plagemann, T. Preuss, and H. König. Architecture, implementation, and evaluation of ConfMan: Integrated WWW and DBS Support for Conference Organization. Technical report, I-1998.016-R, UniK, University of Oslo, Norway, 1998.

[5] T. Kohonen. *Self-organizing Maps*, volume 30 of *Springer Series in Information Sciences*. Springer Verlag, Berlin, Heidelberg, 1995.

[6] R. Mayer, D. Merkl, and A. Rauber. Mnemonic SOMs: Recognizable Shapes for Self-Organizing Maps. In *Proc. of the 15th Workshop on Self-Organizing Maps (WSOM'05)*, 2005.

[7] OpenConf Conference Management System. [Online]. Available: `http://www.openconf.org/`

[8] M. Papagelis, D. Plexousakis, and P. N. Nikolaou. CONFIOUS: Managing the Electronic Submission and Reviewing Process of Scientific Conferences. In *Proc. of the 6th International Conference on Web Information Systems Engineering (Industrial Track)*, New York, USA, 2005.

[9] P. Rigaux. An iterative rating method: application to web-based conference management. In *Proc. of the 2004 ACM symposium on Applied computing*, New York, NY, USA, 2004. ACM Press.

[10] The CyberChair Software. [Online]. Available: `http://cyberchair.org/`

[11] The Microsoft Conference Management Toolkit. [Online]. Available: `http://msrcmt.research.microsoft.com/cmt/`

[12] The MyReview System. [Online]. Available: `http://myreview.lri.fr/`

[13] D. Yarowsky and R. Florian. Taking the load off the conference chairs: towards a digital paper-routing assistant. In *Proc. of the 1999 Joint SIGDAT Conference on Empirical Methods in NLP and Very-Large Corpora.*, 1999.