Intelligent Methodologies for Scientific Conference Management

Marenglen Biba, Stefano Ferilli, Nicola Di Mauro, Teresa M.A. Basile

Dipartimento di Informatica, Università degli Studi di Bari,
via E. Orabona 4, 70126 Bari, Italia
{biba, ferilli, ndm, basile}@di.uniba.it

Abstract. This paper presents the advantage that knowledge-intensive activities, such as Scientific Conference Management, can take by the exploitation of expert components in the key tasks. Typically, in this domain the task of scheduling the activities and resources or the assignment of reviewers to papers is performed manually leading therefore to time-consuming procedures with high degree of inconsistency due to many parameters and constraints to be considered. The proposed systems, evaluated on real conference datasets, show good results compared to manual scheduling and assignment, in terms of both accuracy and reduction of runtime.

1 Introduction

Managing scientific conferences involves many complex tasks regarding the organization and scheduling of the resources involved. Setting up a large conference is often quite a hard task because of the many subtasks involved. This process begins with the conference creation from the Program Committee Chair (PCC for short) who provides essential elements for the conference such as: title, topics, place, start date and deadline of papers submission, start date and deadline of papers review, date of main event etc. In particular, an important activity is the selection of Program Committee Members (PCMs). After this phase the *Call for papers* is issued and the papers submission has to be managed. Then, the next task is assigning each paper to the proper reviewers. During the review process the PCMs evaluate each paper assigned to them and produce their assessment. According to such evaluations the papers are selected and finally approved, which yields an input to the next task for the local organization of the conference. Finally, the authors of accepted papers may submit a new version of their paper, in the so-called camera-ready format, to the PCC, who will send them, together with the preface and the table of contents of the book, to the publisher in order to have the proceedings printed.
Thus, Web-based applications have been developed in order to make easier some of the jobs to be carried out. Typically, these systems provide support for: submission of abstracts and papers by Authors, download of papers by the Program Committee (PC), submission of reviews by the PCMs, handling of reviewers preferences and bidding, Web-based assignment of papers to PCMs for review, review progress

tracking, Web-based PC meeting, notification of acceptance/rejection, sending notification e-mails. However, the more knowledge-intensive tasks require a more "clever" support than just a comfortable interface and improved interaction. Two of the hardest and most time-consuming activities are the assignment of reviewers to submitted papers and the sessions' organization and scheduling (commonly said *Timetabling*). This phase concerns the distribution of the submitted papers into a set of sessions and the assignment of the sessions to available rooms and periods of time. Due to the many constraints to be satisfied, performing manually such tasks is very difficult and usually does not guarantee optimal results.

This work aims at showing how the use of intelligent methodologies can bring important improvements in these two activities, by embedding expert components in a more general framework of a scientific Conference Management System. Section 2 introduces the main features of DOMINUS (DOcument Management INtelligent Universal System), a prototype of a general Conference Management System (CMS), and describes how the expert systems are embedded in it. Section 3 describes the review assignment component, while Section 4 deals with the session scheduling one. Finally, Section 5 concludes the paper, with remarks and future work proposals.

2 DOMINUS for Conference

DOMINUS for Conference is a prototype of CMS. It is still under development. Among other basic services, a CMS can be seen as a system collecting documents (submitted) in electronic form. The distinguishing feature of DOMINUS for Conference lies in its ability to understand the semantics of the document components and content, and to exploit such an understanding to support conference organizers in carrying out their most difficult, knowledge-intensive and time-consuming activities. DOMINUS for Conference is the tailoring to scientific conference management of the general-purpose system DOMINUS for intelligent document analysis, processing and management [7].

DOMINUS implements artificial intelligence methods and techniques for document analysis and understanding. It performs layout analysis on electronic documents in order to exploit the components identified in the document structure to understand their significance and extract from them relevant information about the document. The final aim is to build, through the use of machine learning methods, intelligent semantic-based search engines that offer access to information based on semantic contents of analyzed documents. The entire process of document processing passes through several phases: Document Layout Analysis, that identifies the geometric component in terms of blocks of the document; Document Image Classification that identifies the layout class the document belongs to; Document Image Understanding that identifies the semantic role (according to the document class) of the various layout component identified in the first phase and labels them accordingly; Text Extraction that reads and records the text inside the significant components identified and labeled in the previous phase; Text Categorization during which the text extracted is used to generate key words that describe the topic the document is about; and lastly Information extraction that aims at exploiting all the relevant results of the previous phases to extract and structure significant and important information about the document content.

Once an author has submitted his own paper in an electronic format, it undergoes the processing steps described above. A first exploitation of automatic document processing is that the text, extracted as a result of the six phases above, can be used to automatically check the paper compliance to the conference standards (e.g., layout style and page number) and file the submission record, without requiring the author to manually fill the submission fields (title, authors, affiliations, abstract, number of pages, topics). Then, some of these data (e.g., authors, affiliation and topics) can be exploited by the expert system GRAPE (Global Review Assignment Processing Engine) as an input to carry out the assignment of reviewers to papers. Also the reviewer expertise topics can be automatically inferred by processing their CVs and a selection of their papers: indeed, experience proves that topics manually entered by the reviewers sometimes do not completely represent their actual research interests according to their publications. Lastly, information on paper authors, title and topics can be exploited also by the expert system for session scheduling, in order to carry out its task.

Thus, DOMINUS for Conference aims at defining a multifunction framework for conference management activities that spreads from basic browsing services to scheduling sessions and presentations leading therefore to a complete solution for conference handling based entirely on intelligent components. Its Web version helps organizers of conferences in managing efficiently all their activities through Web accessibility in order to provide a valid tool for conference organizing committees.

3 The GRAPE System

The review process starts after the deadline for paper submission. Then, suitable PCMs are selected, which will act as reviewers, in order to evaluate the submitted papers. Therefore, the PCC sends to individual reviewers the collected submissions with review forms consisting of a set of questions to assess the quality of the paper, that the Reviewers must fill in and return to the PCC. Each submission is typically examined and evaluated by 2 or 3 reviewers. Generally, the review process ends with the PC meeting, where papers' rejection or acceptance for presentation at the conference are discussed on the basis of collected reviews. After this meeting, extracts of the reviewers' comments are typically sent back to all the authors, so that they can improve their paper.

In the current practice, before the submission phase starts, the Chair usually sets up a list of research topics of interest for the conference and asks each reviewer to specify which of them correspond to their main areas of expertise. Then, during the submission process, authors are asked to explicitly state which topics apply to their papers. Such an information provides a first guideline for associating reviewers to papers. However the presence of many constraints makes the problem a very hard one to cope with in terms of time of assignment, for a number of reasons. For instance, some reviewers cannot evaluate specific papers due to conflicts of interest or to their little expertise in the paper topics, or because they are not willing to revise more than a certain number of papers.

GRAPE [1] is an expert system, written in CLIPS (C Language Integrated Production System), that performs the reviewers assignment by taking advantage from both the papers content (topics) and the reviewers preferences (biddings). It can be set to exploit the papers topics only, or both the paper topics and the reviewers' biddings. Its fundamental assumption is to prefer the topics matching approach over the reviewers' biddings one, based on the idea that they give to assignments more reliability. Then, reviewers' preferences can be used for tuning the paper assignments: they are very useful because of the unpredictability of the distribution of reviewers and papers over the list of topics, which causes situations in which some papers have a lot of experts willing to review them, while some others simply do not have enough reviewers.

Let $\{p_i\}_{i=1,\ldots,n}$ denote the set of papers submitted to the conference, t be the number of conference topics, and $\{r_j\}_{j=1,\ldots,m}$ be the set of reviewers. The goal is to assign the papers to reviewers, such that: 1) each paper is assigned to exactly k reviewers (usually, 3 or 4); 2) each reviewer has roughly the same number of papers to review (on average nk/m, but additional constraints can be set to indicate that some reviewer can review at most a given number of papers); 3) papers are reviewed by domain experts; and, 4) reviewers revise articles based on their expertise and preferences.

Two measures were defined to guide the system during the search for the best solutions. The *reviewer's gratification* degree of a reviewer is based on: a) the confidence degree between him and the assigned articles: the confidence degree between a paper and a reviewer is equal to the number of topics in common; and b) the number of assigned papers that the reviewer chose to revise (discussed in more detail in Section 4.2). The *article's coverage* represents the coverage degree of an article after the assignments. It is based on: a) the confidence degree between the article and the assigned reviewers (the same as before); and b) the expertise degree of the assigned reviewers, represented by the number of topics in common. The expertise level of a reviewer is equal to the ratio of topics he is expert in on the whole set of conference topics. GRAPE tries to maximize both reviewer gratification and article coverage degree during the assignment process, in order to fulfill requirements 3) and 4). To reach this goal it is mandatory that each reviewer provides at least one topic of preference, otherwise the article coverage degree would be null.

The two main inputs to the system are the aforementioned sets of papers and reviewers. Each paper is described by its title, author(s), affiliation(s) of the author(s) and topics Tp_i. On the other hand, each reviewer is described by his name, affiliation and topics of interest Tr_j . Furthermore, the system can take as input a set of constraints indicating (i) the possibly specified maximum number of reviews per Reviewer and (ii) the papers that must be reviewed by a reviewer because of specific indication by the PCC. It can be also provided with a set of conflicts indicating which reviewers cannot evaluate specific papers under suggestion of the PCC: in any case, such a set is enriched by GRAPE by deducting additional conflicts between papers p_i and reviewers r_j whenever r_j is a (co-)author of p_i, or the affiliation of r_j is among the affiliation(s) reported in p_i.

We say that a reviewer can revise a paper with a degree that is equal to the confidence degree between them if it is not null, or else is equal to the the expertise degree of the reviewer – ranging in [0,1[– when the confidence degree is 0. Given a paper, the number of candidate reviewers for it is the number of reviewers that can revise it with degree greater than or equal to 1.

3.1 The Assignment Process

The assignment process is carried out in two phases. In the former, the system progressively assigns reviewers to papers with the lowest number of candidate reviewers, thus assuring to assign a reviewer to papers with only one candidate reviewer. At the same time, the system prefers assigning papers to reviewers with few assignments, to avoid having reviewers with zero or few assigned papers. Hence, this phase can be viewed as a search for review assignments by keeping low the average number of reviews per reviewer and maximizing the coverage degree of the papers. In the latter phase, the remaining assignments are chosen by considering first the confidence level and then the expertise level of the reviewers. In particular, given a paper which has not been assigned *k* reviewers yet, GRAPE tries to assign it a reviewer with a high confidence level between them. In case it is not possible, it assigns a reviewer with a high level of expertise.

The assignments resulting from the base process are presented to each reviewer, that receives the list *A* of the *h* assigned papers, followed by the list *A′* of the remaining ones (both, *A* and *A′* are sorted by article's coverage degree). The papers are presented to the reviewer as virtually bid: the first *h*/2 papers of *A* are tagged "high" (he would like to review), the next *h* papers are tagged "medium" (he feels competent to review) and all the others are tagged "low" (he does not want / does not feel competent to review). Now, the reviewer can actually bid the papers by changing their tag, but preserving the number of high and medium ones. Furthermore, he can tag *h*/2 papers as "no". All the others are assumed to be bid as low. Only papers actually bid by reviewers generate a preference constraint. When all the reviewers have bid their papers, GRAPE searches for a new solution that takes into account these biddings. In particular, it tries to change previous assignments in order to maximize both article's coverage and reviewer's gratification. By taking the article's coverage high, GRAPE tries to assign the same number of papers bid with the same class to each reviewer. Then, the solution is presented to the PCC. It is important to say that, if the PCC does not like the solution, he can change some assignments and force the system to give another solution satisfying these new constraints as well. In particular, he may: assign a reviewer to a different paper; assign a paper to a different reviewer; remove a paper assignment; or remove a reviewer assignment. The main advantage of GRAPE relies in the fact that it is a rule-based system. Hence, it is very easy to add new rules in order to change/improve its behavior, and it is possible to describe background knowledge, such as further constraints or conflicts, in a natural way. For example, one could add a rule that expresses the preference to assign a reviewer to the articles in which he is cited.

3.2 Evaluation

The system has been evaluated on real-world datasets built by using data from previous European and International conferences. In order to have an insight on the quality of the results, in the following we present some interesting characteristics of the assignments suggested by GRAPE.

For a previous European Conference the experiment consisted in a set of 383 papers to be distributed among 72 Reviewers, with k = 3 reviews per paper. The system was able to correctly assign 3 reviewers to each paper in 152 seconds. Obtaining a manual solution took about 10 hours of manual work from the 4 Program Chairs of that conference. Each reviewer was assigned 14.93 papers on average by topic (when there was confidence degree greater than 1 between the reviewer and the paper), and only 1.03 papers on average by expertise degree (which is a very encouraging result). GRAPE made many good assignments: in particular, it assigned to 55 reviewers all 16 papers by topics.

Another dataset was obtained from the IEA/AIE 2005 conference. It consisted of a set of 266 papers to be distributed among 60 Reviewers. The conference covered 34 topics. In solving the problem, the system was able to correctly assign 2 reviewers to each paper in 79.89 seconds. The system was also able to assign papers to reviewers by considering the topics only (it never assigned a paper by expertise). In particular, it assigned 10 papers to 38 reviewers, 9 to 4 reviewers, 8 to 6 reviewers, 7 to 1 reviewer, 6 to 9 reviewers, 5 to 1 reviewer, and 2 to 1 reviewer, by considering specific constraints for some reviewers that explicitly requested to revise few papers.

As a conclusion, the rule-based system GRAPE, purposely designed to solve the problem of reviewer assignments for scientific conference management, having been tested on real-world conference dataset, revealed how the use of intelligent methodologies based on expert systems helped in improving significantly the entire assignment process in terms of quality and user-satisfaction of the assignments, and of reduction in execution time with respect to that taken by humans to perform the same process.

4 The Resource Scheduling Expert System

Scheduling activities under rather complex constraints have been subject of much research effort. Indeed, a large number of variants of the problem have been proposed in the literature, that differ from each other because of the type of institution involved and the type of constraints. This problem, that has been traditionally considered in the operational research field, has recently been tackled with techniques belonging also to artificial intelligence (e.g. genetic algorithms, tabu search, simulated annealing, and constraint satisfaction) [2], [3], [4], [5]. The intrinsic complexity of the problem makes it important enough to consider specific knowledge-based solutions in order to provide valid tools that could face the problem bringing significant improvements with respect to manual approaches. Since in most organizations the problem of scheduling is quite overwhelming there are really strong motivations to consider applying artificial intelligence methodologies in order to build expert systems that should deal with the task of allocating resources to activities. In many domains assigning available resources to concurrent activities is critical in terms of importance for the organization processes. Because of the many constraints to be satisfied and many parameters to be considered, manual procedures often yield low quality results even in spite of long execution time spent for the entire process of scheduling.

In the current practice all the activities necessary for the generation of a Timetable are done manually. These activities must deal with the satisfaction of rather complex constraints, such as: unavailability of an author, of the Chair of a session or the unavailability of places necessary for the sessions; pre-assignments that impose certain meetings to be kept at a precise time and/or place; preferences for places or time of a presentation; limitation on place and time, when some places or periods of time are not available for a certain presentation or session; limitation on the sequence and alignment of the sessions when two sessions dealing with the same topic cannot be scheduled to run simultaneously; limitation on the sequence of sessions when the same author cannot attend two simultaneous sessions; reasonable timetable in order to satisfy certain requirements like the maximal number of sessions during a day or the necessity of breaks during lunch time.

4.1 Scheduling in Scientific Conference Organization

One of the activities that involves planning and timetabling is the organization of scientific conferences. The phase of sessions organization requires much effort because of the many parameters to consider. In particular this process consists in the distribution of the papers accepted into sessions and for each session the allocation of available time slices and vacant accommodations. This scheduling activity is heavily constrained by the following variables to be taken into account [6]: Persons: each participant has his own role in the conference thus the scheduling procedure should resolve cases of multiple roles by a person. For instance, an author of a presentation could be Chair of another session, therefore the two sessions cannot be scheduled to run simultaneously; Discussions: these are defined as being any kind of interaction inside the conference like the presentation of a publication, a workshop about a particular topic etc. The constraint the discussions impose on the scheduling process concerns their distribution and organization in non-simultaneous sessions with the same topic and their uniform grouping into session inside which they share the same topics; Sessions: these are the main entities to elaborate in order to produce an optimal timetabling. Constraints directly related to sessions regard the full slice of time necessary for the entire session, the use of specific equipment for the sessions presentations that could be required by another simultaneous session, the presence of not available time periods dedicated to other activities such as lunch breaks which should be taken into account and finally the available accommodations to be shared among concurrent sessions; Timeblocks: in this case the constraints to satisfy are that all the entire time available must be divided in timeblocks which should each contain a session and all the timeblocks should be optimally combined in order to maximally exploit remaining time slices for other sessions; Accommodations: their constraints are related to space requirements of specific sessions, for instance some sessions could be too rich in terms of contents and therefore in terms of participation and need large rooms to proceed. Other sessions could use only particular accommodations because of the particular equipment the rooms provide. Furthermore no accommodations should be shared by different sessions during the same timeblock.

In addition to the constraints introduced above there are different kinds of sessions which means that other parameters must be taken into account. Some of these are: Plenary Sessions that aim at synthesizing the state of the art of a certain research field and that usually run without overlapping with other sessions; Special Sessions that are organized by a person designated on purpose by the Program Committee. These sessions render unavailable the chosen person for other sessions or activities; Sessions of Free Contribution are usually organized by the Program Committee and concern topics proposed by authors; Poster Sessions are organized to allow the authors to present their publications in particular sessions set up on purpose and that differ from the classic session format.

4.2 The Expert System

In order to provide a solid solution for the scheduling of the activities in scientific conference management contexts that could manage to resolve all the constraints presented in the previous paragraph, an expert system was developed and tested in this domain. The previous approaches in resolving this kind of problem have always referred to other solutions already built for other contexts such as university and school timetabling, enterprise activities scheduling, without considering the possibility to design a tailored solution for the context of conference management.

The expert system built is a rule-based system developed in CLIPS. For the construction of the knowledge base of the system there were formulated 11 functions and 30 rules.

Some of these functions were designed to verify the availability of certain resources for a certain session while others verify if the constraints are still satisfied after changing certain parameters. The controls are performed upon all the entities that constitute the scheduling process. We mention here some of the verification tasks the functions implement: one function verifies if the equipment required for a certain presentation is provided by the accommodation to be assigned to the presentation; another function verifies if an author of a presentation who has expressed his unavailability for a certain period can be available for the time that is to be assigned to his presentation; one function verifies if the author of a presentation which is to be assigned a time slice could have to discuss at the same time another presentation in case of more than one paper submitted; another function verifies the parallelism of the session to which is being assigned a certain presentation and other sessions scheduled to be run at the same time; one function calculates the total time of a session.

The rules encapsulated all the human expertise of the scheduling process trying to encode all the peculiarities of such a process. Some of these rules implement actions such as: help the user indicate the accommodation and the time to be inserted for a presentation that must not be parallelized with any other presentation; find the hours available in a room in a certain period of time; eliminate main sessions for which do not exist presentations; break into the availability of an accommodation distinguishing among morning and afternoon hours; change session being based on compatibility between presentation and session topics in case the main session a presentation belongs to, has as estimated time less than one hour and there are not principal session with the same topic; assign a mobile resource to a presentation that requests the resource guaranteeing the fulfilment of all the constraints; resolve the problem that arises when the available remaining time in an accommodation is smaller than the necessary quantity of time to allocate a certain presentation. In this case the rule resolves the problem searching for another presentation on the same topic in the same room which has been assigned a time slice equal to that needed to end the session time.

In order to help the end user to easily manage the process of session management, a stand-alone application was developed in the programming language Java. It offers various visual tools with the purpose of providing to the user user-friendly windows for creating new timetables, modify existing timetables, insert new data, visualize existing timetables, cancel existing timetables, update timetables and visualize details. Through the inserting of new data the user could remove or add new constraints about a presentation or add new data like another accommodation, author or a presentation. Once these data were entered, the expert system running in the background is involved and a new timetable is generated taking into account the new data inserted.

From an architectural point of view the application has 3 levels: presentation level which permits to the end-user to interact with the system through a number of views generated by the system; domain level which implements all the domain-specific functionalities in order to support the interaction with the knowledge base; data source level which is responsible for accessing the knowledge base and is the only component to know the logic and physical structure of the knowledge base.

4.3 Evaluation

The system described in the previous paragraph was tested upon a dataset of a real-world conference, IEA/AIE 2005. The presentations and the person who presented them were 127. Before running the system some basic information was entered like the presentations that could not run simultaneously or eventual breaks between sessions. The execution time of the expert system for building a configuration that satisfied all the constraints was about 4 hours which seems a very enthusiastic result compared to the 7 working days of 4 persons that were fully occupied for creating the timetable. The assignment generated by the expert system maintained a very good level of quality. Through the assignment rules it was possible to allocate automatically about 90% of the presentations. The remaining 10 % was managed using the rules that deal with the incompatibilities and that manage to modify the topic of a presentation based on the assignments made previously.

In addition to the good results in terms of quality and time for the specific task, one should also consider the typical advantages coming from the use of expert systems, such as the exploitation of heuristics to guide the search for a solution, the ability to explain the rationale underlying the response and the possibility of straightforwardly extending the knowledge base with new information, even when hundreds of rules are used.

.

5 Conclusions and Future Work

Organizing scientific conferences involves many complex tasks, whose management requires a more “clever” support than just a comfortable interface and improved interaction. This work showed how the use of intelligent methodologies can bring important improvements in such a domain. Indeed, the scientific Conference Management System DOMINUS for Conference exploits Machine Learning techniques and Expert Systems technology to automatically process submitted documents and effectively/efficiently support conference organizers in two of the hardest and most time-consuming activities: the assignment of reviewers to submitted papers and the sessions' organization and scheduling.

Although results on real-world datasets are encouraging, future work will concern extension and improvement of the system, both as regards its document processing features and as concerns the range of activities supported given by expert components.

References

1. N. Di Mauro, T.M.A. Basile, and S. Ferilli. GRAPE: An Expert Review Assignment Component for Scientific Conference Management Systems. In F. Esposito and M. Ali, editors, Innovations in Applied Artificial Intelligence (IEA/AIE05), volume 3533, LNAI, pp. 789–798, Springer Verlag, 2005.
2. Schaerf A. A survey of automated timetabling. Artificial Intelligence Review, 13(2):87-127, 1999
3. Burke E. K. and Petrovic S. Recent research directions in automated timetabling. European Journal of Operational Research, 140: 266--280, 2002.
4. Carter M.W. A Survey of Practical Applications of Examination Timetabling. Operations Research 34:193-202. 1986.
5. de Werra D. An introduction to timetabling, European Journal of Operations Research, 19:151-162, 1985.
6. Sampson S. E. Practical Implications of Preference-Based Conference Scheduling, Production and Operations Management, 13 (3), pp. 205-215, Production and Operations Management Society, Baltimore, October, 2004
7. Esposito F., Ferilli S., Basile T.M.A., Di Mauro N. Semantic-Based Access to Digital Document Databases. In M.-S. Hacid, Z. W. Ras, and S. Tsumoto, editors, Foundations of Intelligent Systems (ISMIS05), volume 3488, LNAI, pp. 373–381, Springer Verlag, 2005.