
***Cobi*: Communitysourcing Large-Scale Conference Scheduling**

Haoqi Zhang

MIT CSAIL
Cambridge, MA 02139 USA
hqz@mit.edu

Paul André

HCI Institute, CMU
Pittsburgh, PA 15213 USA
pandre@cs.cmu.edu

Lydia Chilton

University of Washington
Seattle, WA 98195 USA
hmslydia@cs.washington.edu

Juho Kim

MIT CSAIL
Cambridge, MA 02139 USA
juhokim@mit.edu

Steven P. Dow

HCI Institute, CMU
Pittsburgh, PA 15213 USA
spdown@cs.cmu.edu

Robert C. Miller

MIT CSAIL
Cambridge, MA 02139 USA
rcm@mit.edu

Wendy Mackay

INRIA
91405 Orsay, France
mackay@lri.fr

Michel Beaudouin-Lafon

Université Paris-Sud
91405 Orsay, France
mbl@lri.fr

Abstract

Creating a good schedule for a large conference such as CHI requires taking into account the preferences and constraints of organizers, authors, and attendees. Traditionally, the onus of planning is placed entirely on the organizers and involves only a few individuals. *Cobi* presents an alternative approach to conference scheduling that engages the entire community to take active roles in the planning process. The Cobi system consists of a collection of crowdsourcing applications that elicit preferences and constraints from the community, and software that enable organizers and other community members to take informed actions toward improving the schedule based on collected information. We are currently piloting Cobi as part of the CHI 2013 planning process.

Author Keywords

Human Computation; Collaborative Planning; Crowdsourcing; CHI; Clustering; Conference Scheduling

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

General Terms

Design, Human Factors

Copyright is held by the author/owner(s).

CHI'13, April 27 – May 2, 2013, Paris, France.

ACM 978-1-XXXX-XXXX-X/XX/XX.

Cobi draws on the crowds in the CHI community

Associate Chairs are 220 committee members who are experts in their area of human-computer interaction. Cobi draws on their expertise to identify groups of relevant papers and proposes to them sessions that they may wish to chair.

Authors of accepted papers are 1000-2000 people who know their own papers well and wish for them to be grouped with, and not opposed to, related papers. Cobi asks them to identify papers that are most related to their paper and those that they would like to see, so as to avoid potential conflicts.

Attendees include 2500-3000 members of the CHI community. Cobi collects their preferences and constraints over papers and sessions in the program schedule. Cobi then engages chairs, organizers, and/or attendees to plan the program based on the information collected.

Introduction

Creating a good program for a large conference is a difficult task. While scheduling hundreds of accepted submissions into sessions across multiple days and rooms, organizers need to consider the multi-faceted preferences and constraints of organizers, authors, and attendees. For example, organizers aim to create sessions of related papers, ensure speakers are not scheduled to give two different talks at the same time, avoid scheduling related sessions opposite each other, and generally keep the program interesting throughout the day for different groups of participants.

We report on the current schedule-creation process at CHI, which involves two main stages. First, once papers are accepted, the technical program chairs and 15-25 associate chairs create small groups of papers, form categories or personas, put together 80-minute long sessions, and build a rough preliminary schedule (see Figure 1). This process is tangible, collaborative, time-consuming, and highly dependent upon the individuals organizing the papers. In the second stage, the conference chairs refine this rough schedule to create the final program. They attempt to resolve conflicts, fix sessions with stray papers that don't fit, and generally look for ways to improve the program. The chairs use a script to check that no presenter is scheduled to be in two places at once, but otherwise make all changes via manual inspection. Past chairs found the process to be very time-consuming. They also noted that resolving conflicts was "painstaking" due to the complexity of the schedule and the lack of feedback on whether changes were resolving existing conflicts or creating new ones.

Despite organizers' best intentions and effort, previous CHI programs often still contain incoherent sessions,



Figure 1. A group of 15-25 people create a preliminary program in person following the CHI PC meeting.

parallel sessions with similar content, and author-specific conflicts. A number of challenges contribute to such problems. First, due to the organic nature of how the committee makes connections between papers in the first stage, many sessions have odd papers mixed in. Second, because the process does not capture the affinities between papers that go beyond the sessions, it's difficult for chairs to make scheduling changes while maintaining cohesive sessions. Third, the committee may not know some of the authors' and attendees' preferences, which can for example lead to sessions of interest being scheduled at the same time. Finally, the lack of software for managing constraints and the sheer size of the schedule make it difficult for chairs to make informed decisions when finalizing the schedule.

Cobi addresses these challenges by drawing on the people and expertise within the community in the planning process (see sidebar). *Cobi* consists of a collection of crowdsourcing applications that elicit helpful information from chairs, authors, and attendees, and software that enables organizers and community members to take informed actions toward improving the program. *Cobi* is currently being piloted for planning CHI 2013, and includes tools for clustering

The Scheduling Problem and the Cobi Solution

Clustering: grouping papers by topic, either into sessions or into affinity groups that are larger than sessions. Cobi recruits community members to group papers in their area of expertise.

Preferences: collecting soft and hard constraints other than topic clustering. By reaching out to organizers, authors, and attendees, Cobi captures individuals' interests and constraints as well as their perspectives on different aspects of the program.

Scheduling: assigning sessions to rooms and time slots, while satisfying hard and soft constraints. Cobi provides a scheduling interface that surfaces conflicts in the schedule and proposes changes to resolve such conflicts.

Session Chairs: assigning chairs to session slots. Cobi proposes sessions for one to chair by automatically computing the best matches to a chair's expertise.

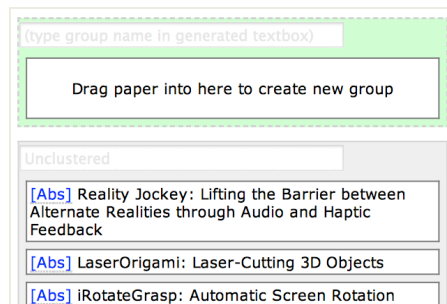


Figure 2. One clustering method presents papers for a contributor to explicitly group into sets of related papers.

papers, collecting preferences, scheduling, and assigning session chairs (see sidebar). By engaging the community in the planning process, Cobi makes the preferences and constraints of its members visible and the planning process more transparent. In doing so, it also shifts the responsibility for the conference program from a few organizers to the entire community at large.

We wish for the entire community to engage with the Cobi system at the Interactivity event to re-plan CHI 2013. In what follows, we briefly describe Cobi's tools for clustering, preference collection, and scheduling.

Clustering

Cobi seeks to better understand the affinities between papers so that similar papers can be grouped together and not placed in opposing sessions. This allows more coherent sessions to be created and attendees to miss fewer sessions of interest. Since automated methods cannot perform this task perfectly and non-expert crowds may only be able to generate broad groupings,

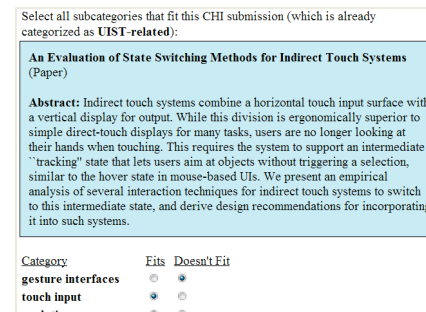


Figure 3. Another clustering method asks contributors to generate category names for papers and to determine whether the categories are a good fit for a paper (shown).

Cobi recruits CHI's associate chairs to group papers in their area of expertise. We explore two alternative interfaces for grouping papers. In one, contributors create groups explicitly (see Figure 2). In the other, contributors first propose category names and then determine which papers fit into which categories (see Figure 3). To promote contributions from many individuals, we keep each set of tasks to ten minutes and display top contributors on leaderboards.

Authors of accepted papers are in a unique position to judge whether other papers are related to their paper. We hypothesize that they also have an incentive to provide input so their paper appears in a session with related papers. Cobi will be presenting authors with papers that are most likely to be similar to their own (currently ranked by TF-IDF) and asking them to judge whether these papers would fit well in a session with their paper. This process helps to collect additional fine-grained affinity information among papers that is useful for session creation and later for scheduling.

The goal is to resolve the remaining conflicts in the schedule and to schedule the unscheduled sessions.

View modes and filtering options provide a multi-dimensional view of the schedule. They allow a user to view helpful context for detecting issues and when making scheduling decisions.

Special Interest Group for

-8 ■ -5

■ -3

Using a constraint solver, the system computes the best slots for a move or swap and highlights them in orange. In this example, swapping with this session would resolve 8 conflicts. Of these conflicts, 5 of them involve authors whose papers are scheduled in simultaneous sessions. 3 involve sessions on similar topics that are scheduled in simultaneous sessions.

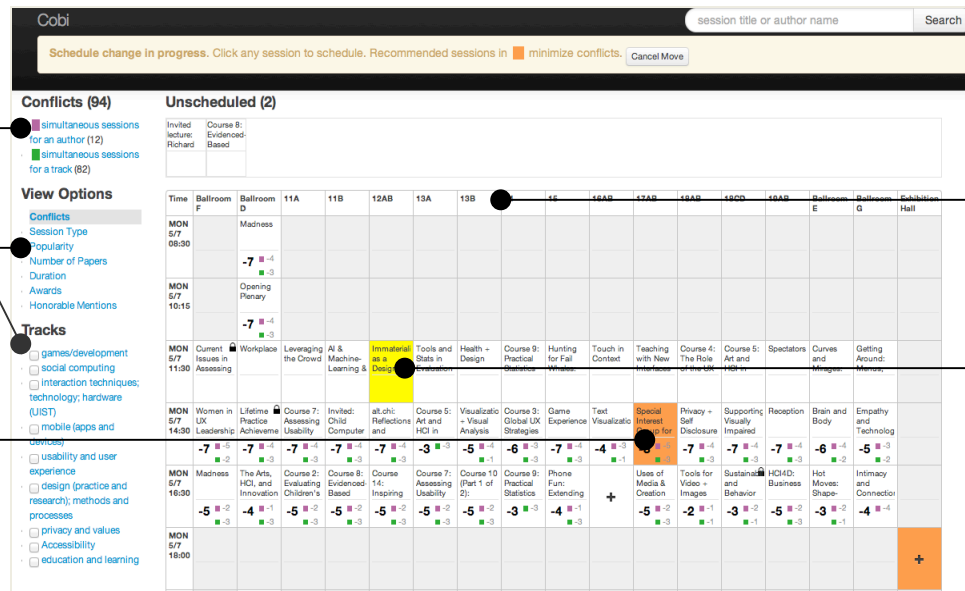


Figure 4. Cobi’s scheduling software displays the effect of swapping the selected session with another session in the schedule in terms of the number of constraints that would be added or removed if the swap were to be made.

Preferences

Cobi allows community members to express their preferences and constraints using natural language and a rich domain-specific language. For example, an organizer can specify times at which a session should take place, identify papers that should be scheduled apart, and note which rooms to assign popular sessions to. We are currently exploring various interfaces for eliciting preferences and constraints from organizers, authors, and attendees for use during scheduling.

Scheduling

Given a set of sessions and collected constraints, Cobi’s scheduling software assists in finalizing the schedule by

The schedule table allows sessions to be directly manipulated, and supports operations such as moving, scheduling, unscheduling, swapping, and locking.

Clicking on a scheduled session and pressing ‘Propose Swap’ (not shown), the system enters swap mode (shown) and displays the effect of swapping the session with other sessions in the schedule.

automatically detecting violated constraints and proposing swaps for resolving them (see Figure 4). The conference chairs will be using the tool to plan the CHI 2013 schedule. As a pilot, we will be releasing another version of the tool that will allow the entire CHI community to help construct and improve the schedule.

Acknowledgements

We thank Patrick Baudisch for helping us pilot Cobi at CHI 2013, Pedro Lopes and Nirmal Patel for sharing their tools and providing technical assistance, and members of the CHI community for contributing to planning CHI 2013 using Cobi.