MIT Open Access Articles

*Frenzy: Collaborative data organization
for creating conference sessions*

The MIT Faculty has made this article openly available. **Please share**
how this access benefits you. Your story matters.

Citation: Lydia B. Chilton, Juho Kim, Paul Andre, Felicia Cordeiro, James A. Landay, Daniel S. Weld, Steven P. Dow, Robert C. Miller, and Haoqi Zhang. 2014. Frenzy: collaborative data organization for creating conference sessions. In Proceedings of the 32nd annual ACM conference on Human factors in computing systems (CHI '14). ACM, New York, NY, USA, 1255-1264.

As Published: <http://dx.doi.org/10.1145/2556288.2557375>

Publisher: Association for Computing Machinery (ACM)

Persistent URL: <http://hdl.handle.net/1721.1/90408>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike

Frenzy: Collaborative Data Organization for Creating Conference Sessions

Lydia B. Chilton¹ Juho Kim² Paul André³ Felicia Cordeiro¹ James A. Landay¹
Daniel S. Weld¹ Steven P. Dow³ Robert C. Miller² Haoqi Zhang^{2,4}

¹University of Washington
Seattle, WA
{hmslydia, felicia0,
landay, weld}@cs.uw.edu

²MIT CSAIL
Cambridge, MA,
{juhokim, rcml}
@mit.edu

³HCI Institute, CMU
Pittsburgh, PA
{pandre, spdow}
@cs.cmu.edu

⁴Northwestern
University
Evanston, IL
hq@northwestern.edu

ABSTRACT

Organizing conference sessions around themes improves the experience for attendees. However, the session creation process can be difficult and time-consuming due to the amount of expertise and effort required to consider alternative paper groupings. We present a collaborative web application called *Frenzy* to draw on the efforts and knowledge of an entire program committee. *Frenzy* comprises (a) interfaces to support large numbers of experts working collectively to create sessions, and (b) a two-stage process that decomposes the session-creation problem into meta-data elicitation and global constraint satisfaction. Meta-data elicitation involves a large group of experts working simultaneously, while global constraint satisfaction involves a smaller group that uses the meta-data to form sessions.

We evaluated *Frenzy* with 48 people during a deployment at the CSCW 2014 program committee meeting. The session making process was much faster than the traditional process, taking 88 minutes instead of a full day. We found that meta-data elicitation was useful for session creation. Moreover, the sessions created by *Frenzy* were the basis of the CSCW 2014 schedule.

Author Keywords

Crowdsourcing; groupware; communitysourcing

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI):
User Interfaces – Graphical user interfaces.

INTRODUCTION

When planning an academic conference, organizers group papers into thematic sessions so that attendees can see related talks in the same time-block. However, constructing

and arranging conference sessions can be a challenge, especially for a small group of organizers. First, organizers often need to consider a large number of accepted papers from multiple sub-disciplines. Knowledge about these papers is typically distributed within the community. Second, accepted papers can be grouped in multiple ways (e.g., by topic, by problem domain, by method of study), so maintaining a global outlook of how and why papers fit together is often non-trivial. Third, sessions must satisfy two hard constraints: each paper must be assigned to exactly one session and every session should be the same length. This implies that even coherent paper groupings may be infeasible if they contain too many or too few papers, and that clever alternative groupings may be required to avoid stray papers or incomplete sessions. Since creating a session affects what other sessions can be created, the process of coming up with coherent sessions that satisfy global scheduling constraints requires effective coordination.

To manage this process, many conference-organizing committees use printouts of the accepted paper abstracts. Through informal observations and interviews with organizers of two large conferences, we learned that a small group of dedicated organizers typically spend a day or two in person creating sessions by printing abstracts on cards, then tangibly arranging cards in piles, and trading cards between piles, until the piles roughly form session-sized groups. After all this, the information on the cards is manually entered into a computer.

This time-consuming process has several shortcomings. First, the one-to-one correspondence of a paper to a physical card limits the number of people that can work on grouping a paper at one time. Second, connections between papers are often made organically as an organizer walks around the room to collect related cards. This can lead to sessions with odd papers mixed in. Third, the time constraints and difficulty of navigating through piles of cards makes it difficult to consider parallel alternatives for grouping papers. Organizers often feel "locked in" with these initially created sessions, because any modification requires them to add to or break apart another session. Further, to account for stray papers, organizers often leave the meeting

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CHI 2014, April 26–May 1, 2014, Toronto, Ontario, Canada.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-2473-1/14/04...\$15.00.

<http://dx.doi.org/10.1145/2556288.2557375>

with thematic sessions that contain too many or too few papers, leaving additional work for refining the schedule.

We introduce *Frenzy*, an alternative approach for creating sessions that draws on the distributed knowledge of the entire program committee. First, *Frenzy* breaks the task of session making into two sub-problems: *meta-data elicitation* and *global constraint satisfaction*. During meta-data elicitation user can see all the accepted papers, search them by text, and add two types of meta-data: they can add suggested categories for a paper and they can indicate that they like a suggestion category and think it has “high session making potential” by clicking a “+1” button. This stage allows us to use many experts in parallel to contribute their intuitions about likely sessions. This stage can be done during the breaks of the program committee meeting, or after the wrap-up meeting. The next stage is global constraint satisfaction which uses a smaller group of volunteers who are co-located and can communicate easily with one another to use the suggested categories and “+1” votes to solve the global constraint of assigning all papers to exactly one session that has 3 or 4 papers.

To enable both stages we introduce a web application called *Frenzy* that facilitates parallel collaboration among large and small crowds. *Frenzy* uses some familiar concepts from social media such as Twitter’s “feed” and Facebook’s “tagging” and Google Plus’ “+1” button to help view and add meta-data to papers. It also has standard features such as search and autocomplete. *Frenzy* allows parallel collaboration by providing each user with their own view of the data which they can search without affecting other users, but which propagates all their meta-data to other users immediately to eliminate redundant meta-data and give an active sense of collaboration.

We evaluate *Frenzy* by deploying the tool to the CSCW 2014 Program Committee, at which 48 committee members contributed to *Frenzy*. The session-creation meeting took 88 minutes compared to the traditional process, which often takes all day.

This paper makes the following contributions:

1. We introduce *Frenzy*, a collaborative web interface for constrained data organization that uses **goal completion with actionable feedback** to alert users of what needs work, while providing them freedom to choose their task.
2. We address the challenge of using groups at different scales to collectively satisfy constraints and achieve a cohesive global structure by **decomposing the problem of creating sessions into two sub-problems**: meta-data elicitation and constraint satisfaction.
3. **We evaluate *Frenzy* by deploying it to a conference program committee over two days.** We show that providing actionable feedback allowed users to pursue their own strategies for completing the goals. We show that the meta-data elicited was useful for session creation.

Moreover, the meta-data collected from *Frenzy* was more useful than the legacy paper categories used by that conference.

The paper proceeds as follows. We first describe related work in crowdsourcing and groupware. Next, we discuss the design motivation for *Frenzy* and describe the system and its implementation. We chronicle the deployment of *Frenzy* to the CSCW 2014 Program Committee (PC) with data analysis and interviews. The sessions created by *Frenzy* were the basis of the CSCW 2014 schedule.

RELATED WORK

Frenzy is a collaborative tool that builds on results from groupware, crowdsourcing, communitysourcing and design.

Groupware has a long history in HCI and has proven to be unexpectedly challenging. In analyzing several expensive failures of groupware systems, Grudin [9] observes that developing for groups is more difficult than developing for individuals and he supplies eight challenges of developing groupware systems. His insights on social processes are at the core of *Frenzy*’s design.

The existing solution for creating conference sessions is a social process. There are people in a room reading over accepted papers, accumulating knowledge of some of the papers, and talking to people with other knowledge hoping to form connections between papers and dealing with conflicts and failures as they arise. As Grudin describes, social process are often guided by personalities, tradition and convention. Computers don’t have access to this knowledge and thus our challenge is to provide support to the process without seeking to replace these behaviors. Grudin also points out that especially for social systems, there is often a big difference between our description of how the system work and how the system actually works. Social systems have many exceptions to the rule, unforeseeable errors and aspects of personality that we don’t know how to account for. The challenge for *Frenzy* is to remain flexible and not impose rigid structures that will fall apart as exception and errors naturally arise.

Crowdsourcing takes a different approach to collaboration than groupware. Crowdsourcing – particularly microtask based crowdsourcing – avoids many of the social complications of groupware by replacing the current social process with its own workflows. Crowdsourced workflows [3,12,13,16] stitch together the results of microtasks which workers perform in isolation. A benefit of workflows is that computers understand and can optimize them [7]. However, the downside is that they can be perceived by users as inflexible and not enabling true collaboration, which might not matter on platforms such as Mechanical Turk, but does matter when the users are in a community with shared goals and knowledge which can be leveraged. *Frenzy* was influenced by the notion of workflows. Instead of having a detailed microtask-based workflow, it simply

has two stages, each one with freedom as to what data the users want to see and what contributions they want to make.

Microtasks have the benefits that users know what to do, they know that it will take only a few minutes, and they don't have to worry about the big picture. These factors tend to make contribution more attractive and more easily parallelizable. Successfully breaking a problem into microtasks is challenging, but there are common microtasks have emerged in several crowdsourcing systems. Voting is a common microtask used in [14, 16, 20]. Soliciting tags, categories, or labels has been used in the ESP Game [1] and many other systems [4,6,8]. Frenzy builds on the idea of microtasks by allowing users to make contributions as simple as tagging and voting, but integrating these contributions into an interface that is open and collaborative rather than combining them in a workflow.

Frenzy addresses the challenge of facilitating collaboration: How do you create a single artifact, such as a conference program, out of many diverse opinions? Most crowdsourcing systems avoid this challenge by picking problems that parallelize well, such as labeling images. However, a few crowdsourcing systems do output a single artifact. Often machine learning or other statistical techniques are used to create artifacts such as transcribed text [15], a clustering of data [2, 8, 20] or a taxonomy [6].

A departure from the statistical approach is a crowdsourcing trip-planning tool called Mobi [21]. Trip-planning is a constraint satisfaction problem, and Mobi solves it by presenting a no-workflow interface where users can choose their own contribution from a list of "to do" items. Frenzy builds on ideas in Mobi by introducing a platform for parallel work, a two-stage workflow, and by generalizing the notion of "to-do" items into "actionable feedback," where the feedback can be changed based on the stage of the system and what goals need the contributors need to focus on.

Another general approach which is introduced in the design literature is "flare and focus" also called "divergence and convergence" [5]. This is a technique used in ideation processes such as brainstorming where many ideas are solicited in the "flare" stage and then a few of them are chosen to be deepened in the "focus" stage. This was the inspiration for Frenzy's two-stage collaboration process.

Frenzy is part of a growing literature on communitysourcing [10] and conference planning [2, 11, 19]. Frenzy builds on a previous conference scheduling research tool called Cobi [11] which allows the conference organizers to schedule sessions at times that abide by constraints such as "no author can have his papers scheduled for presentation at the same time." Cobi starts with a preliminary set of unsched-

uled papers based on paper sessions grouping. It allows the organizers to select times for each session, and possible swap papers in between session to avoid conflicts. Frenzy is a tool that helps create the initial sessions and affinity data if papers need to be swapped.

SYSTEM DESCRIPTION

The Frenzy interface (Figure 1) consists of four sections: the *query bar*, *results feed*, *results hyperbar*, and goals with *actionable feedback*.

The *query bar* is a good place for users to start exploring Frenzy by performing a text search over papers. A paper matches the query text if the paper's title, author list, author affiliations, or abstract contains that text. Text search helps users narrow the list of papers by broad terms such as "Crowdsourcing", and also helps retrieve particular papers by an author or from a keyword in the title. Users can also see all the papers by clicking the "Show all papers" button.

The results of the query are displayed in the *results feed* which shows a vertical list of *data cards* (Figure 2) for each of the papers returned by the query. A *data card* displays the papers information (title, authors, and abstract) on the left side and the user-generated *meta-data box* on the right. The meta-data box contains a list of categories suggested for the paper, as well as a count of how many people have added a "+1" vote to the category, indicating that they think they category has high session potential. If the users are in the constraint satisfaction stage of Frenzy, there is also a text box to enter a session name. Users may add their own meta-data by adding a category, adding their own "+1" vote, adding or editing the session name (if applicable) or removing a category from the paper. We call this area the results feed because it draws some similarity to the social media concept of a feed of updating information. The interface draws visual connections to the Twitter interface, which is familiar to most of our users.

When the results feed is updated by a query, Frenzy also update the *results hyperbar*. The results hyperbar displays feedback about the query that was performed – how many results are returned, and a statement of the query. It also returns additional filters which the user can apply to their query. For example, if the user searches for "Crowdsourcing" the results hyperbar will say "26 papers in 'Crowdsourcing'" and will list up to 5 categories the user can filter by, such as 'paid crowdsourcing.' If the user clicks these additional filter categories, the results will update to have "5 Results for papers in 'Crowdsourcing' and 'paid crowdsourcing.'" These additional filters make it easy to drill down into large categories.

Figure 1. Frenzy interface, highlighting 4 sections: actionable feedback, query bar, results hyperbar, and results feed.

The left panel of Frenzy displays the *actionable feedback*. In both stages of Frenzy (meta-data elicitation and constraint satisfaction) there are two goals for users to achieve as a group. For example, in the meta-data elicitation stage one of the goals is: "Every category must have at least two papers in it (No singleton categories.)" Instead of instructing users how to achieve this goal, Frenzy provides two types of feedback on progress towards that goal that users can easily take action on (hence the name actionable feedback.) One type of actionable feedback for this goal is the list of categories, with the number of papers in that category in parenthesis. Any categories with only one paper are displayed in red text, indicating a problem. Users can then click on that category to see what paper is in it, and either remove the paper from it (thus deleting the category), or

add more papers to the category. An additional type of actionable feedback is the number of papers that meet the goal and the number that do not. By clicking the number of papers that do not meet the goal, users generate a query that returns papers that need work. This provides an easy and direct way to find papers in need of contribution.

Frenzy is a web app implemented in node.js and Bootstrap. Each user logs in and then sees the Frenzy interface. When users query the system, their query is private to them. However, all meta-data generated is propagated to other users within 5 seconds. If a user adds a category to the paper, that category will appear in other users' data cards, and in the actionable feedback pane, the number of papers in that category will update, as well as the number of completed goals, if applicable.

Figure 2. The data card for paper cscw663. The left side is paper details and the right side is the meta-data box.

To make Frenzy a flexible microtask platform, our design has three goals:

1. Enable lightweight contributions.
2. Allow users to access all the data and tasks.
3. Promote completion of goals with actionable feedback tailored to user groups and stages of problem solving.

We next describe how the Frenzy interface supports these design goals.

Design Goal 1: Enable Lightweight Contributions

In order to encourage as much participation as possible even in only short periods of free time, all contributions to Frenzy are small tasks that a user can complete in a under a minute. Users can choose their contributions and move easily between tasks. This allows users to make contributions that best fit their availability and expertise.

All user contributions are made using the meta-data boxes associated with each meta-data card. There are four ways a user can make contributions on each meta-data card:

Add a category. If a user can think of a new category the paper fits into, they are encouraged to add it. Autocomplete in the “add a category” textbox helps users reuse categories that are already in Frenzy. Additionally, categories with only one paper appear in red to indicate that they are singletons.

Remove an existing category. In order to remove a paper from a category, the user simply unchecks the category. Unchecked categories become less visually salient by turning grey, and their upvote button disappears. Category names remain visible and can be checked again to reassign a paper to a category.

Upvote a category. We expect every paper to be assigned to multiple categories. Categories may represent different aspects of a paper, such as its topic (e.g., education, question answering, games, health) or contribution type (e.g., study, system) Categories may also vary in their levels of

granularity or quality. Up-voting a category indicates that the category has high potential of becoming a session. Typically this means the category is small (3-8 papers) rather than being overly broad such as “Empirical methods” which has over half the papers in it.

Create/update/delete session names. Users can place a paper in a session by entering a session name in the meta-data box. Session assignments can be deleted or edited at will. To distinguish sessions from categories, Frenzy only allows each paper to be placed in a single session at any given time. This feature is only turned on during the constraint satisfaction stage.

Giving users more control and freedom also comes with potential disadvantages. First, users have to navigate the set of tasks and take the time to figure out how they will contribute. Second, seeing current work could potentially bias the results of future work. While some of these issues are mitigated by the actionable feedback presented to users, they also represent some inherent tradeoffs of having no fixed microtask workflow.

Design Goal 2: Give users access to all data and tasks

Lightweight contributions are convenient, but are only useful if users can find a place to contribute. Fixed workflows present users with a designated place to work, but we rejected the idea of using workflows in Frenzy because we wanted users to make contributions based on their expertise. Nobody knows a user’s expertise as well as the user does, thus our solution to enabling contributions was to offer access to all papers and tasks through domain-specific search features. Because papers are attached to tasks to be done on those papers, users can search for papers that fit their expertise as a way for searching for tasks to be done that match their expertise. Frenzy supports text search over the titles, authors, affiliations and abstracts of the papers. Frenzy enables category-based searches in the actionable feedback panel, and the results hyperbar allows search results to be further filtered by relevant categories.

A benefit of this search-based solution to giving access to tasks is that it supports the existing discussion-based social process rather than replacing it. Users can discuss papers and then search for them, see their details and add meta-data that may results from that conversation. Search provides the user control and freedom [17] that communities want to feel in their collaborative efforts.

Design Goal 3: Promote completion of goals with actionable feedback

Showing hundreds of items gives users the control and freedom that they want, but can also be overwhelming. Frenzy provides actionable feedback to allow users to quickly find places to work that need their attention and which they are knowledgeable about.

In order to effectively coordinate work between a large group of users providing meta-data and a small group of users making sessions, we associate with each sub-problem its own set of goals, based on which to present feedback.

In the case of meta-data elicitation, we set two goals:

MDE1. Every category must have at least two papers in it (No singleton categories)

MDE2. Every paper needs to be in a least one category with +1 for session-making potential

In the case of session constraint satisfaction to assign all papers to session, we set these two goals:

SCS1. Every paper needs to be in a session

SCS2. Every session must have more than 2 papers.

Associated with every goal is visible feedback. For example, the MDE2 goal is displayed at the top of the actionable feedback panel with two buttons, one saying “100 Items Complete” and the other saying “35 items need work.” This feedback is actionable because the user can click on the button and filter the results to only the items that need work, or that are already completed. For the MDE1 goal, Frenzy highlights any singleton categories in red. When the user clicks on the red category name and sees which paper it contains, they can either remove the singleton category from it (which deletes the category from the system because it then will have no papers in it), or find other papers for the category.

It is important to note that Frenzy does not assign tasks to users. The affordance in the actionable feedback is a shortcut for a search that nudges users towards finding a subset of the data to attend to. Once users find a place to work, there are many strategies for meeting the goals. A user may find that they need to add categories to a particular item, or remove them from another, or merge two categories, or split large categories into multiple smaller categories. Results from the psychology literature indicate that setting goals is an effective management style that gives people freedom in how they choose to achieve a goal [18].

DEPLOYMENT

We partnered with the chairs of the CSCW 2014 Program Committee (PC) and deployed Frenzy for the initial session creation process. The traditional process involves 10-15 PC members meeting face-to-face after all the paper decisions have been made with the accepted papers’ information printed on cards that they organize into piles which then become sessions. This process has two drawbacks: 1) it can only involve a limited number of people, and therefore fails to leverage the expertise of all the members of the communities, and 2) The amount of exploration is limited by that fact that there is only one copy of each paper that must be assigned to exactly one pile at any particular time. This process tends to take the better part of a day.

Traditionally, session making can only start after all the paper decisions are final because the problem has constraints that don’t make sense to solve until the data is complete. In contrast, Frenzy breaks down the problem into two stages: meta-data elicitation (referred to hereafter as MDE Frenzy) and session constraint satisfaction where the session-making is finalized (referred to hereafter as SCS Frenzy). MDE Frenzy and SCS Frenzy use the same Frenzy platform, but with minor adjustments to the actionable feedback portion of the user interface. In MDE Frenzy, the goals are MDE1 and MDE2 (see previous section) and the actionable feedback focuses on categories. In SCS Frenzy, the goals are SCS1 and SCS2 and there is actionable feedback for both categories and sessions. All the meta-data gathered in MDE Frenzy appears in SCS Frenzy.

The CSCW PC received over 500 submissions. Of those, approximately 100 were fast-tracked for likely acceptance and approximately 100 were slated for discussion. We loaded the ~100 fast-tracked papers into Frenzy before the meeting started, and as the committee made decisions about individual papers, they were manually added to the Frenzy interface. This way, Frenzy always contained a current view of the accepted papers.

CSCW has a set of 32 *legacy categories* for papers such as “Entertainment/games” and “Social Network Site Design and Use.” The authors of the paper selected multiple of these categories that apply to their paper. We imported this pre-existing meta-data into Frenzy as categories. As a result, all papers had at least two categories at the start of Frenzy.

Meta-Data Elicitation (MDE) Frenzy

The PC meeting had 63 attendees, 43 of whom participate in meta-data elicitation for Frenzy. At PC meetings, there are several times when certain members must step out of the room and into the hallway due to conflicts with the authors of the paper being discussed. We used this free time by setting up Frenzy on a computer with a large monitor in the hallway. Thus, PC members could browse the accepted papers and enter meta-data during free time. Since PC members tend not to take their laptops with them into the

hallway, this was the only computer available, and often multiple people gathered around the screen and discussed the papers and meta-data together. PC members also used the interface during breaks from the meeting and from their own laptops inside the meeting. After all the paper decisions were complete, the PC members engaged in a 5-minute session dedicated to entering meta-data in Frenzy.

Session Constraint Satisfaction (SCS) Frenzy

After the PC meeting ended, nine volunteers including the PC chairs moved to a smaller location to create a preliminary set of sessions for the conference. Sessions needed to have between 3-5 papers in them, ideally four. The number did not need to be exactly four since sessions were likely to change when refining the schedule, e.g., if presenters have conflicts. The goal was to have initial sessions to work from, and the actionable feedback encouraged sessions to have at least 3 papers.

The group started with the meta-data collected from the all the PC members, including 330 category entries made and 236 category upvotes.

DATA ANALYSIS

During the two days of Frenzy deployment, we logged all user actions: sign-ins, sign-outs, queries, and data entries. We observed people using the system and conducted interviews with the SCS Frenzy participants. We now analyze this data to evaluate Frenzy. We show that the overall goal of session making was achieved in record time and successfully incorporated the view of a large group of experts. We show that the design of breaking the problem into two sub-problems was effective by showing that the meta-data collected in MDE Frenzy was used extensively in SCS Frenzy. We show that goals with actionable feedback in both MDE Frenzy and SCS Frenzy provided the users the control and freedom to define their own strategies for successfully achieving their goals.

Achieving overall goals

Frenzy saw substantial usage over the 2-day deployment. A total of 48 participants contributed over 10.28 hours of usage. In that time, 2,365 queries were issued and 1,088 meta-data contributions were made. Over 250 contributions were made during the 5-minute period where all PC members were asked to spend 5-minutes simultaneously using Frenzy. This demonstrates the capacity of Frenzy to allow simultaneous contributions from a large group of users.

During SCS Frenzy usage, the 9 volunteers to complete the sessions achieved the goal of assigning every paper to a session and having every session have 3-5 papers in 88 minutes.

Breaking Session-Making into 2 Sub-problems

Frenzy breaks session-making into 2 sub-problems: MDE Frenzy and SCS Frenzy. We evaluate this design by testing whether the meta-data from MDE Frenzy was useful for

making sessions in SCS Frenzy. The usefulness of category meta-data was tested by looking at how many of the papers ended up in a session that matched one of the categories give for it. The usefulness of upvote meta-data was tested using a logistic regression to model the effect of upvoted categories on session creation. We defined a category as “matching” a session if one or more important keywords were shared between the category name and the session name.

How useful was category meta-data?

In total, 93 of 135 papers (68.9%) had a category that matched a session. The categories on those 93 papers could have come from two sources: the 32 predefined categories assigned by the PC or the 63 non-singleton categories contributed by users of MDE Frenzy. Although both are helpful in making sessions, only categories contributed by users of MDE Frenzy can be claimed as a benefit of asking users to add meta-data. Of the 93 papers with matching sessions, 40 of the matches came from predefined categories and 53 came from user contributions. MDE Frenzy more than doubled the number of useful categories.

How useful was upvote meta-data?

MDE Frenzy generated 99 non-singleton categories. After constraint satisfaction, there were 34 sessions, of which 25 matched categories. We want to know if the +1 voting for categories helped to determine which of the categories would be turned into sessions. To test if +1 voting provides a useful signal that a category will match a session, we run a logistic regression predicting the probability that a category will match a session ($prMatchesSession$). The dependent variable is an indicator of whether that category was the most +1 upvoted category for at least one paper ($wasMostUpvoted$).

$$prMatchesSession \sim a + b * wasMostUpvoted$$

Coeff.	Estimate	Std. Error	Pr(> z)
a	-2.0794	0.4330	1.57e-06
b	1.7658	0.5278	0.000821

Table 1. A logistic regression predicting whether the probability a category matches a session is dependent on whether that category was the most upvoted category at least once.

Table 1 shows the coefficient estimates. A chi-squared test shows coefficient b is statistically significant ($p = 0.0082$). The interpretation of this logistic regression model is that for a category where $wasMostUpvoted = 0$, the predicted probability that it will match a session is 0.12. In contrast, for a category that was the most upvoted category for at least one paper ($wasMostUpvoted = 1$), the predicted probability is 0.42. Thus, if a category is ever the most upvoted category for a paper, it has high session-making potential, which confirms the usefulness of upvote meta-data to session-making.

Goals and Actionable Feedback in MDE Frenzy

The main mechanism that drives workers to make progress in Frenzy is having goals with actionable feedback. This helps users find a subset of the data on which to work in order to achieve the goals. Because we do not know what actual items are best to give to each worker and we do not know what tasks to give them (adding labels vs. removing labels vs. placing items into sessions), we provide an interface that grants users control and freedom (Design Goal 2) that allows them to find their own strategies for contributing towards the goal (Design Goal 3). In our deployment of MDE Frenzy we found 3 strategies for contribution:

Browse and Edit

Users often came to Frenzy because they were curious to browse the list of accepted papers and to see how the conference was shaping up. As they browsed, the meta-data box was clearly visible, and often they made a contribution. For example, by browsing the system, a PC member who was not an expert in education added the category “MOOCs.” Another user noticed several papers about teens, adolescents or children and added the category “youth” based on several key word searches for “children,” “teens,” and “adolescents”. The categories “email”, “facebook” and “twitter” were added by users who noticed the terms pop up, searched for them to see if they were themes in the program, and then added the categories.

Check for Patterns within One’s Area of Expertise

The papers already had good category labels given by the authors from a checklist provided by the conference organizers. Thus workers could enter Frenzy, select their subfield of expertise and look over the existing meta-data. For example, a crowdsourcing expert selected the “Crowdsourcing” category, which had 26 papers. He looked over the list and found that a few of them had something in common. He added the category “Crowdfunding” to five of the items. These five papers were ultimately grouped together in a

session called “Crowdfunding: Show me the money!”

Similarly, a social media expert searched for “Social Networking Site Design and Use” which had 16 papers and created the category “Politics/Social Media” which had 4 items. These four items were ultimately grouped in session called “Social Media & Politics.”

Direct contribution

As PC members were actively discussing papers in the meeting, they would occasionally see connections or patterns during the meeting then want to come out to the hallway to enter them into the system. For example, during the PC meeting discussions, one PC member realized two papers were accepted about social media and depression (“Social Structure and Depression in TrevorSpace” and “Characterizing and Predicting Postpartum Depression from Facebook Data”). At the next opportunity, they used Frenzy to do a text search for “depression” and added “Depression” as a category to both papers.

Goals and Actionable Feedback in SCS Frenzy

As shown earlier, the session making goals were completed in record time and SCS Frenzy made heavy use of meta-data collected in MDE Frenzy. Figure 3 shows data entry activity averaged over 1-minute intervals. Types of data entry and color codes and stacked. From this graph and the interviews conducted with the users, we identify four distinct stages of the session making process: additional data-entry, removing clutter, session making, and lastly session negotiation.

Stage 1: Additional data entry and low hanging fruit

Figure 3 shows that during the first 28 minutes, more categories were added and more upvotes were contributed. In addition, some of the more obvious sessions were created, such as turning the 4 papers in the category “Crowdsourcing” into a session called “Crowdsourcing: Show me the money!” We call such sessions made directly from categories low hanging fruit. During this period, volunteers largely worked alone without communicating to one another. One volunteer stated:

I started with stuff I knew... I have a pretty good sense if a paper belongs there or not.

Stage 2: Removing Clutter

Figure 3 shows that from minutes 24-32, many categories were removed. A large number of the more obvious sessions had been created but now the low hanging fruit was gone. The volunteers pointed out a number of categories that were created for fun such as “Pacific Northwest Pride” “What’s Niki doing now?” They felt these were cluttering the data and making it hard to find meaningful meta-data, and so they had them removed. One volunteer commented:

I tried to remove some of the joke labels because they were starting to clutter up the user interface.

Figure 3. Data entry activity during SCS Frenzy. Data entries are broken into 4 types and stacked on the graph: category updates, session updates, category removal, and category upvotes.

Stage 3: Session making with discussion

Figure 3 shows that from minutes 32 to 62, more sessions were created. In contrast to the beginning of the process where volunteers with taking low hanging fruit and not communicating, there was lots of communication both to coordinate efforts and to get reassurance from others. Our volunteers said:

A couple of times I would try to make a session, but confirm verbally I was making that session [so] that they didn't have to do it.

People would be working on their stuff, then somebody would speak up and say. Where could this paper possibly go?

Stage 4: Session Negotiation

Figure 3 shows that from minutes 60 to 88, category addition slowed with many gaps in time where no sessions were updated. During this period, there was a lot of discussion about how to resolve the most difficult remaining constraints. Discussion dominated data entry, and the discussion resembled a negotiation with volunteers making proposals of switches and coming up with creative ways of reconfiguring existing sessions. For example, they took the 7 papers in “Q and A”, found an additional paper on social networks and then split “Q and A” into two categories, “Q and A” and “Friendsourcing” which focused on using Facebook and Twitter friends for information needs. This creative problem solving is characteristic of the session negotiation stage.

User feedback

Our interviews revealed a consensus that Frenzy was an improvement over past experience with paper-based conference session-making:

It was definitely much less painful to create those initial sessions, [but that] it was more fun to do it with a social tool... it was nice to do it in a group, to have a few other people in the room ... especially when we ran into tricky ones, helpful to talk it over.

There was also consensus that face-to-face aspects of the collaboration during SCS Frenzy were still vitally important to the process:

I wouldn't want to do this remotely, having everyone there was important, awareness of what other people were doing, touching certain papers, or thinking about sessions.

In particular, the social interactions in the room were essential for resolving some of the session making problems:

Sometimes people would go in different directions and make simultaneous edits. We had to stop and negotiate a little.

DISCUSSION

While Frenzy is a special purpose tool for creating conference sessions, it solves a more general problem: collaborative data organization. There are three techniques Frenzy uses that are critical to enabling the collaborative process within the community.

Technique 1: Flare and focus

As discussed in the related work section, flare and focus is an approach to ideation introduced in the design literature. Frenzy's two-stage process borrows from this idea: meta-data elicitation, wherein users generate categories and contribute +1 votes is similar to the flare stage, and constraint satisfaction, where the main focus is eliminating options and focusing on the best of the generated options. The flare stage is what allows Frenzy to use a large crowd of collaborators in parallel without worrying about how the constraints will be solved and the focus stage is where the constraints are solved.

One important detail of Frenzy's flare and focus process is that there is such a thing as too much flare. Between the flare stage and the focus stage, it is necessary to get rid of clutter created in the flare stage – categories clearly too large or too small, joke categories. Eliminating clutter is crucial to a successful focus phase.

Technique 2: +1 Voting

Voting is a common crowdsourcing technique. +1 Voting is slightly different from true voting because users don't necessarily have all the options set before them when they cast their +1 vote. +1 voting is an analog of voting which is not limited by the synchronization barrier: you don't need to control the workflow of microtasks to ensure that voting is only done after all the options are listed.

A subtle point about +1 voting is that if you want to include pre-existing categories, you need to have a mechanism like +1 voting in order to set reasonable goals for the users. Without pre-existing categories, a simple goal such as “every paper must have at least one category” will suffice. However, if there are pre-existing categories, that simple goal may already be satisfied. Changing the goal to say “every paper must have at least one category with a +1 vote” the means the users will look at every paper, even if it has already been categorized.

Technique 3: Social interaction

Frenzy uses design patterns from crowdsourcing such as tagging and voting microtasks. However, microtasks alone, even in a platform that affords a great deal of user control and freedom, are not enough to solve the problem of session making. Face-to-face communication is crucial for solving many small but important issues: resolving conflicts, coming up with new terms that will be acceptable to the community, and arriving at a consensus that the process is complete and the results, while not perfect, are satisfactory. One of the major lessons Grudin [9] derives from the failures of groupware systems is that social processes have subtle inner workings that we do not always understand, are difficult for computers to have access to and reason about, but are essential to the process. Thus we should not seek to replace the existing social process with a rigid workflow, but to add lightweight microtasks on top of the current process.

Frenzy is a hybrid approach which can take advantage of the lightweight and efficiently parallelized contributions of microtasks and still incorporate the inspiration and social cues that drive consensus which transpires from face-to-face communication. Our general observation is that microtasks do not have to be distributed to strangers in rigid workflows, they can be made social and can enhance face-to-face interaction rather than aim to replace it.

FUTURE WORK

The area of biggest potential improvement in Frenzy is to speed up Stage 3 of Constraint Satisfaction: Session Making with Discussion. Frenzy could be extended to offer machine learning based analysis to suggest paper groupings to decrease the amount of time spent searching for connections after the low-hanging fruit has been taken.

There is also the potential to extend the techniques for collaborative data organization to other domains. One area that has shown early promise is collaborative creation of photo albums. For example, guests of a wedding could all be invited to upload photos they took, browse, categorize and add +1 votes and collectively output an album using each photo at most once, and not more than 200 photos.

CONCLUSION

In this paper we present Frenzy, a tool for collaborative data organization applied to the complex task of conference session making. Frenzy enables groups of experts to simultaneously contribute by breaking the problem into two sub-problems: meta-data elicitation and session constraint satisfaction. Frenzy gives users control and freedom in identifying their own strategies for accomplishing set goals. Actionable feedback promotes meeting goals and steers users toward useful work.

In our deployment of Frenzy at the CSCW 2014 PC meeting, we evaluated the actionable feedback features by identifying three strategies used to satisfy the goals. We showed the benefit of breaking the problem into two sub-problems by showing that the meta-data collected in MDE Frenzy helped form more sessions that the categories created by the PC before the PC meeting. Moreover, the sessions created by Frenzy were the basis of the CSCW 2014 schedule.

ACKNOWLEDGEMENTS

This work is supported by Quanta Computer as part of the Qmulus project, the Ford-MIT Alliance, NSF awards SOCS-1111124 and SOCS-1208382, the WRF / TJ Cable Professorship, ONR grant N00014-12-1-0211, and NFS Grant IIS 1016713 and IIS-1110965. Juho Kim is partly supported by the Samsung Fellowship. Lydia thanks Kristi Gable, Katlyn Edwards, and Janelle Van Hofwegen.

REFERENCES

1. von Ahn, L., Dabbish, L. Labeling Images with a Computer Game. CHI 2004.
2. Andre, P., et al. Community Clustering: leveraging an academic crowd to form coherent conference session. HCOMP 2013.
3. Bernstein, M., et al.. Soylent: a word processor with a crowd inside. UIST 2010.
4. Bigham, J. P., et al.. VizWiz: nearly real-time answers to visual questions. UIST 2010.
5. Buxton, B. 2007. *Sketching User Experiences: Getting the Design Right and the Right Design*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.
6. Chilton, L. B., Little, G., Edge, D., Weld, D., and Landay, J. Cascade: crowdsourcing taxonomy creation. CHI 2013.
7. Dai, P., Mausam, Weld, D.S. Decision-Theoretic Control of Crowd-Sourced Workflows. AAAI 2010.
8. Gomes, R., Welinder, P., Krause, A., and Perona, P. 2011. Crowd-clustering. In NIPS '11 (2011).
9. Grudin, J. 1994. Groupware and social dynamics: eight challenges for developers. *Commun. ACM* 37, 1 (January 1994)
10. Heimerl, K., Gawalt, B., Chen, K., Parikh, T., and Hartmann, B. CommunitySourcing: engaging local crowds to perform expert work via physical kiosks. CHI 2012.
11. Kim, J., et al. Cobi: A Community-Informed Conference Scheduling Tool. UIST 2013.
12. Kittur, A., Smus, B., Khamkar, S., Kraut, R. E. CrowdForge: Crowdsourcing complex work. UIST 2011.
13. Kulkarni, A., Can, M., Hartmann, B. Collaboratively crowdsourcing workflows with turkomatic. CSCW 2012.
14. Lasecki, W.S., et al. 2013. Chorus: a crowd-powered conversational assistant. UIST 2013.
15. Lasecki, W., Miller, C., Sadilek, A., Abumoussa, A., Borrello, D., Kushalnagar, R., and Bigham, J. Real-time captioning by groups of non-experts. UIST 2012.
16. Little, G., Chilton, L. B., Goldman, M., Miller, R. C. TurKit: human computation algorithms on mechanical turk. UIST'10.
17. Norman, D.A. 2002. *The Design of Everyday Things*. Basic Books, Inc., New York, NY, USA.
18. Pink, D.H. Drive: The surprising truth about what motivates us. Penguin (2011).
19. Sampson, S.E. Practical implications of preference-based conference scheduling. *Production and Operations Management* 13, 3 (2004), 205–215.
20. Tamuz, O., Liu, C., Belongie, S., Shamir, O., and Kalai, A. Adaptively learning the crowd kernel. In ICML 2011.
21. Zhang, H., Law, E., Miller, R. C., Gajos, K. Z., Parkes, D. C., and Horvitz, E. Human computation tasks with global constraints. CHI 2012.