

Crowdboard: Augmenting In-Person Idea Generation with Real-Time Crowds

Salvatore Andolina¹, Hendrik Schneider^{2,3}, Joel Chan⁴, Khalil Klouche²
Giulio Jacucci^{1,2}, Steven Dow⁵

¹Helsinki Institute for Information Technology HIIT,
Department of Computer Science, Aalto University, Finland

²Helsinki Institute for Information Technology HIIT,
Department of Computer Science, University of Helsinki, Finland

³University of Göttingen, Göttingen, Germany

⁴HCII, Carnegie Mellon University, Pittsburgh, PA, United States

⁵Cognitive Science Department, UC San Diego, La Jolla, CA, United States

ABSTRACT

Online crowds can help infuse creativity into the design process, but traditional strategies for leveraging them, such as large-scale ideation platforms, require time and organizational effort in order to obtain results. We propose a new method for crowd-based ideation that simplifies the process by having smaller crowds join in-person ideators during synchronous creative sessions. Our system Crowdboard allows online crowds to provide real-time creative input during early-stage design activities, such as brainstorming or concept mapping. The system enables in-person ideators to develop ideas on a physical or digital whiteboard which is augmented with real-time creative input from online participants who see and hear a live broadcast of the meeting. We validate Crowdboard via two user studies in which dyads of in-person ideators brainstormed with the help of crowd ideators. Our studies suggest that Crowdboard can effectively enhance ongoing brainstorming sessions, but also revealed key challenges for how to better facilitate interactions among in-person and crowd ideators.

ACM Classification Keywords

H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous

Author Keywords

Creativity Support Tools; Brainstorming; Crowdsourcing.

INTRODUCTION

Creativity is a social process. Contrary to the “lone genius” myth, recent studies point out that innovation is more likely to arise through combinations of ideas from collaborating individuals with a diverse set of viewpoints and experiences [26,

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.
C&C '17, June 27–30, 2017, Singapore, Singapore

Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-4403-6/17/06...\$15.00

DOI: <http://dx.doi.org/10.1145/3059454.3059477>

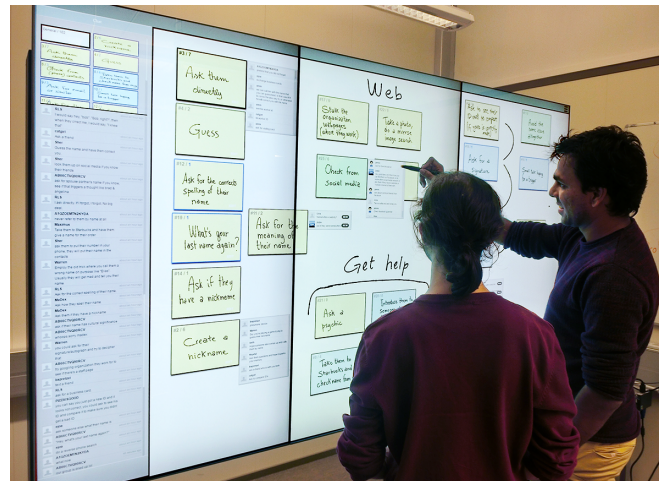


Figure 1. Crowdboard 2.0 (studio interface). In-person ideators write or sketch ideas on virtual sticky notes and get real-time creative input on their ideas from crowd workers who see their activity and who contribute to the creative discussion. Crowd workers can add their ideas either to chats attached to specific notes or to a general chat.

35, 39]. Based on this consideration, researchers have recently investigated a variety of techniques for crowd-based ideation [4, 9, 10, 41]. A number of large-scale ideation platforms have been developed to create spaces where people contribute ideas to help solve challenging problems. The guiding principle is to build up a large pool of ideas from which a few good-quality ideas can be extracted. This asynchronous large-scale approach to crowd-aided ideation has some limitations, however. For example the quality of individual ideas is generally poor [13]. Many obvious ideas are also proposed multiple times, due to the lack of coordination among crowd members [33]. Furthermore, a great amount of time is required to select the best ideas [37].

In-person ideators who may want to leverage the crowd to boost an early-stage idea would be subject to these and other limitations. In particular, they would need a great amount of

time to collect and filter the crowd-generated ideas, before iterating on them. Also, the asynchronous large-scale innovation format does not allow interaction between crowds and those who proposed the creative topic. Because of this, the conversation cannot be shaped by new perspectives opened by idea submissions that have already been collected, and important opportunities may be missed. Although researchers have explored various strategies to address the quality problem and to help crowds produce better ideas in a more efficient way [11, 33, 41], the issues caused by the lack of interaction between crowds and in-person ideators still persist. In this work, we explore this design opportunity and propose a different crowd-based ideation paradigm in which in-person and crowd ideators work together in real-time.

We develop a method by which *in-person ideators can elicit synchronous creative input from dozens of nonexpert crowd workers who are recruited on-the-fly* to help the ideators generate solutions to challenging problems. The proposed approach may benefit both in-person and crowd ideators. The crowd's creative and diverse input would help in-person ideators to avoid missing important opportunities in early design stages where the cost of a large deviation from an initial idea is not prohibitive. One promising strategy for team support during brainstorming is to provide different sources of inspiration through context-aware provocative stimuli [3, 38]. Nonexpert crowds could provide a more effective source of inspiration through a better-contextualized and more creative input. Furthermore, crowds could be able to contribute to the discussion with original ideas. In-person ideators would lead the conversation and act as facilitators, which might help crowds generate more high-quality ideas [11]. Cognitive theories of idea generation suggest that the active exchange of knowledge could be beneficial as it causes mutual stimulation of associations [32]. However actions need to be taken to mitigate the effect of negative social processes [12]. In this context the risk of evaluation apprehension may be lower since designers and online crowd worker are less likely to know each other.

In this work we make the following contributions:

- 1) We develop a method for ideation in which colocated ideators brainstorm together with dozens of nonexpert crowd workers who are recruited on the fly. We embody this idea in the Crowdboard system, a (physical or digital) whiteboard that allows online crowds to see and hear an ongoing brainstorming session and to contribute by sending ideas or inspirations to the workspace that the crowd shares with the people hosting the discussion. We describe the iterative design process of Crowdboard.
- 2) We describe two laboratory experiments to understand how real-time input from nonexpert crowds can support in-person ideators. In Study 1 we informally investigate the feasibility of the Crowdboard idea for supporting early stage design sessions in real-time. We learned that crowd members can understand the conversation and successfully provide their creative input and that in-person ideators value real-time crowd input and incorporate it into their discussions. In Study 2 we go a step further by investigating whether real-time input from nonexpert crowds leads to improved brainstorming productiv-

ity/outcomes. Our results showed that participants generate more creative ideas over time using Crowdboard compared to using the whiteboard without a crowd.

BACKGROUND

Crowdsourcing in the design process

Many have explore the potential of using online crowds to help solve complex [7] and creative problems [21, 22, 24]. For example, Dow et al. [13], explored the use of crowdsourcing in the classroom to contribute at different stages of the innovation process, such as needfinding, ideating, testing, and pitching. CrowdCrit [28] leveraged paid crowdsourcing to generate and visualize high-quality visual design critique. Voyant gave users access to a nonexpert crowd to receive structured feedback on the perceptions of their designs from a selected audience [40]. Crowd vs. Crowd (CvC) [31] used competitions between crowd teams, led by a designer, to crowdsource the entire design process. One limitation of these systems was that they focused on asynchronous interactions and highly structured, iterative, tasks. Other approaches [25] leveraged real-time crowd workers for creating prototypes, but were limited by the lack of communication among involved parties. In contrast in this paper, we explore how creative idea generation tasks can be accomplished in-real time by collaborating teams of in-person and crowd ideators.

Supporting effective idea generation

In-person brainstorming [30] is a popular technique for collaborative idea generation in organizations. A popular practice in brainstorming is to write ideas on sticky notes and apply them on large surfaces such as walls or whiteboards. The contents on the walls can then be easily rearranged, regrouped or replaced as needed. However physical tools have their limitations. For example it is difficult to update physical notes, to keep track of different workspace modifications, and to transfer content to permanent storage for later analysis. To overcome the limitations of physical tools while providing the advantages of large surfaces, many computer-based interventions for brainstorming (and creative collaboration in general) have explored the use of digital whiteboard and tabletop systems [15, 16, 17, 18, 23]. Those systems allow fluid collaboration among ideators and easy interaction with generated ideas, but they often provide little support for cognitive stimulation. Our work extends prior work on digital whiteboard systems by augmenting the workspace used by in-person ideators with novel inspirations and ideas that are generated in real-time by nonexpert crowds who are recruited on-the-fly.

Other computer-mediated interventions provide users with context-aware stimuli to help spark novel associations. For example Idea Expander supports group brainstorming by intelligently selecting pictorial stimuli based on the group's conversation [38]. Similarly InspirationWall [3] uses automatic speech recognition to monitor users' discussions and automatically suggests keywords to support idea generation. Although those approaches attempt to intelligently provide inspirations based on the topic currently being discussed, automated methods have limited intelligence. In contrast, with Crowdboard we leverage the collective intelligence of real people (who

listen to the discussion and contribute with context-aware and creative input) by bringing a level of personalization, improvisation, and even playfulness that is still out of reach for even the best automatic approaches.

Leveraging crowds for ideation

The idea to leverage the wisdom of the crowd [36] in idea generation is not new. The typical approach is to maximize the quantity of generated ideas and to then select a few quality ones [9]. However the quality of individual ideas is generally poor and the selection of the best ones is very costly. For example, for its “10 to the 100th” project, Google had to recruit around 3000 employees for the selection process, pushing the project more than 9 months behind schedule [37]. Those limitations raise doubts about the efficiency of the approach and have prompted researchers to find ways to improve the quality of crowd-generated ideas. IdeaHound intervened by helping crowds find inspirational ideas from a large pool of ideas [33]. Yu et al. [42] used nonexperts to identify distant domains that have the potential to yield useful and non-obvious inspirations for solutions. Those domains are then used to find inspirational examples that can be adapted to solve the original problem. Chan et al. used abstracted solution paths to inspire crowds and studied how various sense-making approaches affect crowd ideation [10]. With IdeaGens experts monitored incoming ideas through a dashboard and offered high-level inspirations to guide ideation. Controlled experiments showed that experienced facilitators increase the quantity and creativity of workers’ ideas more so than unfacilitated workers do [11].

In our work, we abandon the asynchronous large-scale ideation paradigm and the quantity-over-quality approach. We involve smaller crowds but aim to maximize quality rather than quantity by leveraging real-time collaboration. One of the main interventions proposed in this work is to virtually bring crowds into the design studio, and have them join the team of in-person ideators. Unlike other approaches that have a separation between crowd and in-person ideators (who can only interact offline), Crowdboard allows for the active exchange of ideas between parties, during a synchronous idea generation session. Crowd and in-person ideators have complementary strengths. A beneficial attribute of the crowd is the diversity of its expertise and ideas. Providing diverse ideas could inspire the in-person ideators. The crowd also brings the perspective of “outsiders” (not fixated on what has already been done), which complements the in-person group’s context. On the other hand in-person ideators can benefit from face-to-face interactions and more sophisticated tools. They can also lead the discussion, which can improve crowd performance [27]. To our knowledge Crowdboard is the first system that combines those complementary strengths by having crowd and in-person ideators joining forces in the same synchronous creative session. However, synchronous interactions bring their own challenges: in in-person brainstorming, various social processes (such as production blocking, evaluation apprehension, social loafing) can lead to significant productivity losses [12, 43]. While electronic brainstorming systems have mitigated some of these concerns (e.g., reducing production blocking with parallel idea entry [43]), these systems have much smaller group sizes (on the order of 10 or so members),

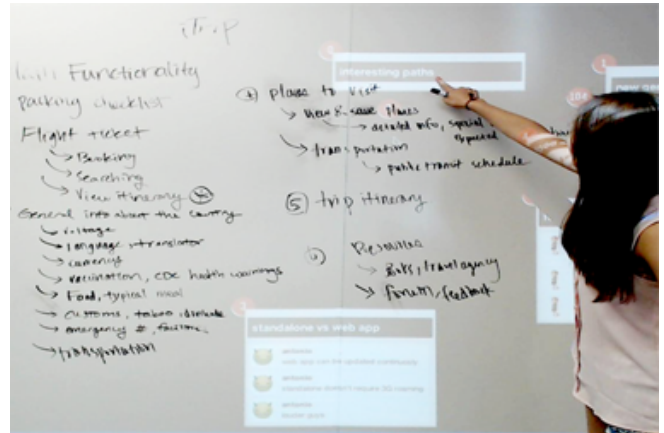


Figure 2. Crowdboard 1.0 (studio interface). A physical whiteboard is augmented with crowd ideators’ comments projected onto the physical board.

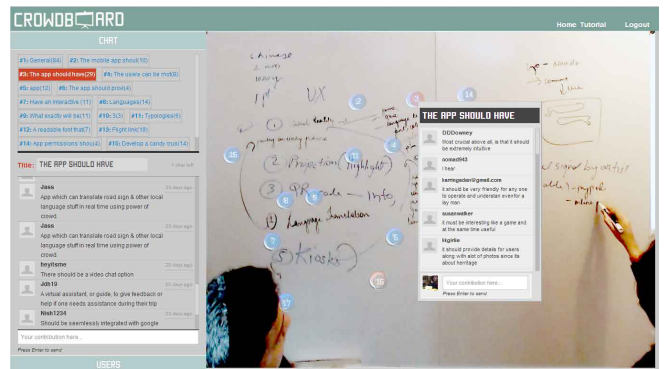


Figure 3. Crowdboard 1.0 (web interface). Crowd ideators see a list of design conversations (upper left) and the whiteboard activity (right). They can leave comments that, in turn, are projected onto the physical board.

who are typically either all in the same room or remote, and typically homogeneous in terms of level of expertise, compared to our envisioned hybrid in-person and crowd teams. The increased team size, heterogeneous expertise, and hybrid in-person/remote dynamics yield unique challenges which we address in our iterative design work.

CROWDBOARD 1.0

Integration with Physical Design Studio

Our first version of Crowdboard allows in-person ideators to leverage creative crowd input while working together in a physical studio. Crowdboard’s in-studio components comprise a traditional whiteboard, a webcam, a projector, a Microsoft Kinect, and a laptop. Real-time interaction is supported by a web-server, which is responsible for storing the system state and synchronizing the clients.

When the system is active, the crowd sees a video stream of the whiteboard and hears audio from the designers. In-person ideators draw on the whiteboard with regular markers and interact with the projected crowd comments using touch gestures (Figure 2). The crowd-generated discussions are positioned on the whiteboard in contextually-appropriate locations, allowing the conversations to specifically refer to the content drawn on

the board. Each discussion thread displayed on the board has a title and one or more comments from the participants. In-person ideators can expand or collapse the discussions using tap gestures and move them around by tapping and dragging the discussion markers.

Interactions among the crowd workers

Crowd workers interact with the team using the web interface (Figure 3). The right panel contains a synchronously updated view of the studio whiteboard. Workers can create a new discussion thread at a particular position on the board by double-clicking. They can also expand an existing thread and then add to the discussion. The left panel contains a list of the current discussion threads, the controls for the live audio broadcast of the meeting, and a list of other workers.

STUDY 1: CROWDBOARD'S FEASIBILITY

In Study 1 we informally investigate the feasibility of the Crowdboard idea for supporting early stage design sessions in real-time. We aim to understand: 1) whether crowd members understand the conversation and successfully provide creative input; 2) whether teams value this real-time crowd input and incorporate it into their discussions.

Participants

Via social media and the university's behavior research participant pool, we recruited five teams of 2-3 design students (12 total) who had a range of design experience. We also recruited 90 online participants from Mechanical Turk. The online participants were mainly Americans (65%) and Indians (31%), and they ranged from 19 to 68 years old (average: 33); 59% were male, and 39% were female. They represented a wide range of occupations (firemen, librarians, managers, etc.), but a majority were in the IT (24%) or education (13%) fields. We tried two approaches from the literature for recruiting crowds that would be available at a predetermined time. For our first 3 trials, we recruited workers in advance (30,30, and 45 workers, respectively), and we asked them to return at a set time [8]. This yielded 8, 11, and 12 workers, respectively. For our fourth and fifth trials, we began recruiting workers immediately before our task so as to have a larger pool [6], and this resulted in 33 and 26 participants respectively.

Procedure

We conducted two long brainstorming sessions of about 35 minutes each, and three short sessions of 15 minutes each. The design brief asked each team to discuss and create a concept map for the idea of traveling to a foreign country. After each session, we asked both the teams and the crowds to fill out a survey. We also interviewed the in-person ideators to get a better understanding of their thoughts about the experience.

Results

In each session, the teams started discussing the system functionality and the design aspects for the problem in question, and then they started generating a concept map. Meanwhile, crowd workers followed the conversation, and added annotations and comments to the virtual whiteboard. In the design studio, the teams noticed the new annotation icons as they

were projected on physical whiteboard, and they saw the issues raised by online participants. The in-person ideators and the crowds continued to work together to fill the whiteboard with comments, ideas, and potential solutions.

Effectiveness of collaboration

We conducted a video analysis of the sessions and coded the following events: 1) An in-person ideator reads a comment aloud; 2) An in-person ideator adds an idea or potential solution to the concept map; and 3) An in-person ideator adds an idea or potential solution to the concept map as a direct consequence of a crowd comment.

The following examples give an idea of how we counted the crowd contributions:

Crowd participant CP18 comments "I love the idea of turning it into a game". In-person participant P4 then reads that comment aloud and adds "Game" to the concept map.

Crowd participant CP11 comments "You can turn it into a game to collect all the experiences...". P4 reads the comment aloud and adds, "Role Playing scavenger hunt" to the concept map near "Game".

Figure 4 shows the number of crowd contributions that were added to the concept map. We observed that Crowdboard greatly affected the design conversation. In our experiments, on average, 35% of the ideas generated in each session came from crowd input (SD = 11%), a rate of 0.56 ideas per minute (SD = 0.23).

Interviews with the in-person ideators who participated in the longer sessions suggest that they considered the crowd input beneficial. They confirmed that crowd participants helped generate ideas: "I was relying a lot on the crowd" [P4]; "I think the conversation was good... People all had good opinions" [P1]. In the shorter sessions, however, the in-person and crowd participants generally complained about the lack of time, which suggests that longer sessions are necessary for better discussions. Participants in particular felt overwhelmed by the large number of comments coming in such a short time.

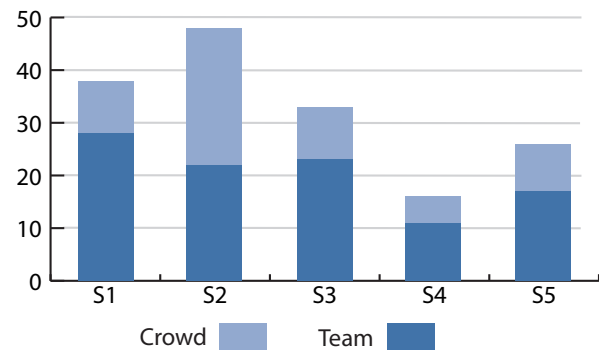


Figure 4. Experiment 1 - CrowdBoard 1.0. Number of ideas and potential solutions added to the concept map. The light blue represents the number of ideas and potential solutions added as a direct consequence of a crowd comment.

Crowd feedback

The post-session surveys asked the crowd members to assess, using a 1–10 scale (with 10 being the best) the extent to which 1) they understood the conversation; 2) they felt the team had acknowledged their input; 3) the team adopted their ideas. The results were 1) $M = 7.84$; $SD = 1.92$; 2) $M = 6.83$; $SD = 2.68$; and 3) $M = 6.53$; $SD = 2.68$; Those results suggest that Crowdboard made the crowd members feel like their voices had been heard.

To get a better understanding of the crowd experience, we conducted a content analysis on the open-ended feedback from the crowd workers and identified 4 main themes. About 36% of workers felt like their voices had been heard, confirming our intuition. The following are examples of comments coded with that theme:

“It felt rewarding to get feedback and see our ideas being valued as potentially important” [CP16]

“I enjoyed hearing everyone’s opinions and felt like everyone’s opinions were valued by those who were hosting the chat...” [CP34]

Another theme that emerged from the analysis was that the system made the crowd workers feel that they were part of the conversation with the in-person ideators. About 70% of the workers mentioned this theme:

“The session was engaging, I felt the team spirit despite the distance.” [CP50]

“I felt like I was very much a part of the conversation as I saw some of the ideas/suggestions I had been addressed.” [CP72]

“It was like being part of the team.” [CP69]

Furthermore, most crowd workers (about 59%) showed enthusiasm about the experience:

“I really enjoyed the experience...I had fun, and felt as if my time was well spent...” [CP76]

“I really enjoyed the experiment. I was a little sad when it ended. I felt like I was impacting others in a positive way.” [CP15]

“Best HIT I’ve done on mTurk. I felt as if I was actually in the room and enjoyed contributing to the conversation.” [CP38]

Although crowd workers reported that the general experience was positive, some workers also reported that they did not feel that they had any impact on the conversation because their input was not acknowledged or because they did not manage to provide novel contributions because of the great amount of ideas that other participants had already generated. About 19% of the workers felt overwhelmed by the excessive number of comments.

Discussion

Perception of crowd input

The results from Study 1 suggest that generic crowds can be successfully engaged in synchronous design discussions with local team members. They also suggest that the teams valued the crowds’ input, with a large percentage of crowd-generated ideas included in the final mindmap.

Reliability of crowd workers

Though online crowds of workers can often include workers who either intentionally or unintentionally provide problematic

input, we observed no clear instances of this. One possible explanation is that this is a result of the visible, synchronous interaction with end users. Additionally, only 20% a workers left the task before it was complete, which is exceptionally low for Mechanical Turk workers given the length of the task. This suggests that Mechanical Turk workers can be reliable for this kind of synchronous experiments.

Design implications

The lessons learned from the study have several design implications:

- *Crowd acknowledgment.* In Study 1 teams that took great care in trying to engage the crowd throughout the entire session, such as by talking with them instead of just talking to each other, had better crowd engagement and participation than the groups that did not do so. This suggests that actions need to be taken to ensure adequate crowd acknowledgment.
- *Limitations of depth-camera sensing.* Lighting conditions can greatly affect the quality of interaction with projected crowd comments. Techniques based on depth-sensing cameras may lack the necessary robustness and could thus lead to poor interaction. Additionally, occlusion problems can occur, making the detection of touch events impossible. In Study 1, we used a “Wizard-Of-Oz” technique to provide smooth interaction with projected elements, as our goal was not to perform a proper test of the system, but rather to investigate the feasibility of the Crowdboard idea. However the touch interaction problem needs to be addressed before systems like Crowdboard can be accepted in the real world.
- *Notifications.* A good notification system is needed to take full advantage of crowd input. In Study 1 the in-person ideators missed many relevant crowd comments. A division of roles between in-person ideators helped mitigate the issue. For example, in the second session one team member decided to continuously monitor the crowd input while the other team member led the discussion and added ideas to the whiteboard. This method helped make the team productive (see Figure 4). However a better notification system is needed. The challenge is making notifications as unobtrusive as possible while maintaining their effectiveness. Using audio or other prominent notifications may be too overwhelming for in-person ideators considering the large amount of events that the system generates.
- *Crowd recruitment, assignment review, and payment.* A time-consuming phase in each Crowdboard session concerns the management of crowd work. To save time and make Crowdboard accessible to users with no experience in crowd management, an integrated system would be needed. This would allow users to manage operations such as recruitment, task review, and payment with just a few clicks from the main interface.
- *Maintaining context consistency.* In Crowdboard 1.0 crowd workers could create new discussion threads at specific X-Y locations on the board so that they could refer to a specific idea. Threads (and associated comments) would in turn be projected in the physical board on the design studio. A limitation of such an approach is that when an idea is

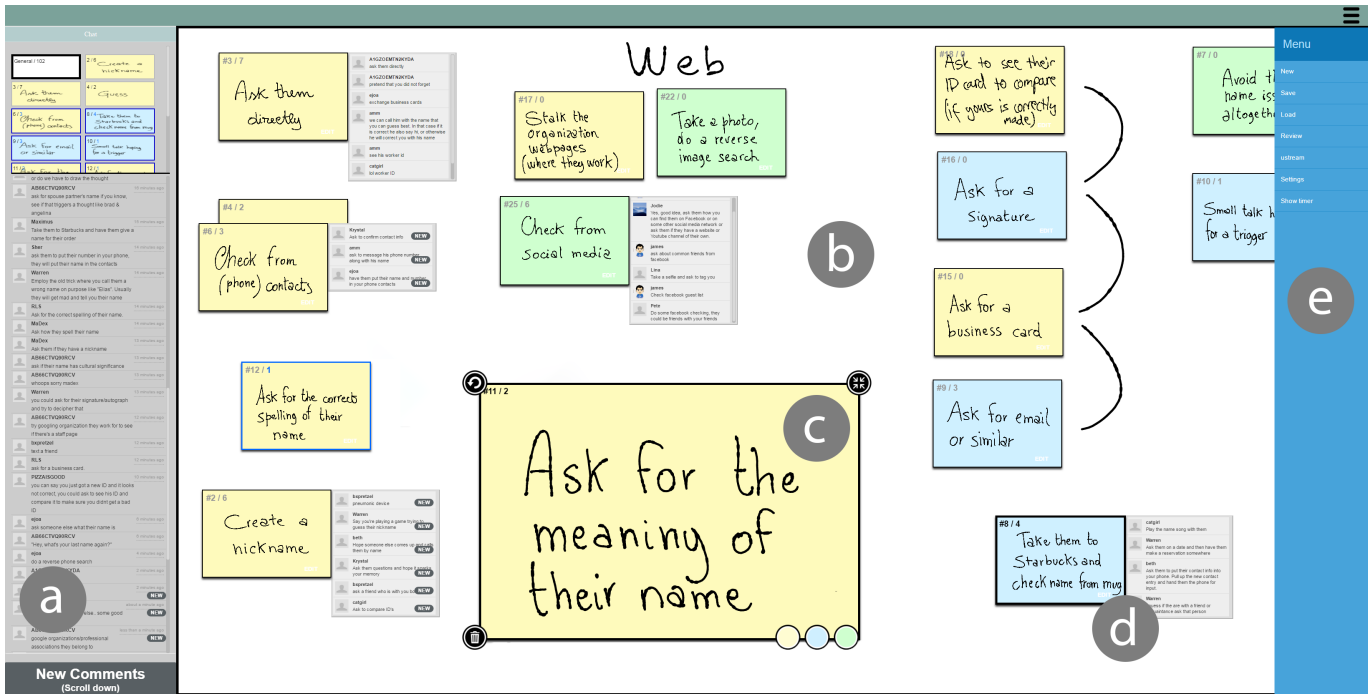


Figure 5. CrowdBoard 2.0 (studio interface). a) Left panel: General chat. b) Main panel: White canvas where users can draw with IR-pens and erase with barcode-based erasers. This area can contain virtual sticky notes, which are created with special barcode-based tools. c) A virtual sticky note in edit mode. d) A virtual sticky note and its associated chat. e) The menu for accessing session- and crowd-management widgets.

erased and added back to another location on the board, the corresponding discussion threads do not follow that action. In such cases the user is responsible for moving the thread accordingly to maintain consistency; this operation can distract from the main task. Future iterations of the interface need to take into account the need to maintain context consistency.

CROWDBOARD 2.0

Based on the lessons learned in Study 1, we decided to iterate and redesign Crowdboard. The goal was to bring Crowdboard to its full potential so that real-time crowds' effects on ideation could be formally studied.

A digital whiteboard with a crowd inside

Taking into account the design implications from Study 1 we decided to develop Crowdboard 2.0 using a big touchscreen composed of three 55" modular MultiTaction Cells [29]. The screen uses a vision-based technology that allows the tracking of unlimited touch points, including hands, fingers, fingertips, 2D barcodes, and IR-pens, via the TUIO protocol [20, 29]. This design choice allows in-person ideators to overcome the limitations of depth-camera sensing. One of the main benefits of Crowdboard 1.0 – the possibility to use traditional tools such as regular markers – is also preserved in Crowdboard 2.0, which provides good approximations of the real tools in IR-pens and barcode-based tools (Figure 6). We designed Crowdboard 2.0 to be a stand-alone system that people with no crowdsourcing experience could use. We used the API from Amazon Mechanical Turk to enable crowd recruitment, assignment reviews, and payment from the main interface.

Studio interface

The studio interface is depicted in Figure 5. The interface is divided into two main areas: the left panel (Figure 5a) and the main panel (Figure 5b).

The main panel is a white canvas where users can freely draw and erase using IR-pens and barcode-based erasers (Figure 6). Special barcode-based tools (Figure 6) allow the ideators to create virtual sticky notes. In normal mode, these notes can be freely moved and arranged on the workspace. Tapping on the bottom-right corner of a note brings up edit mode (Figure 5c). When in this mode, the notes expand to allow for easier handwriting and/or sketching. Buttons on the lower part of the note allow the user to change the background color or delete the note. On the top part of the note, the “undo” button can be used to cancel the last action that occurred on the note, and the “shrink” button can be used to save the note and exit edit mode. When a note is created, a new chat is also generated and attached to that note (Figure 5d). The chat serves as a bucket for crowd ideas and comments that build on what is written on the note. The numbers shown on the top




Figure 6. CrowdBoard 2.0 tools: IR-pens for drawing and barcode-based tools for erasing and creating virtual sticky notes.

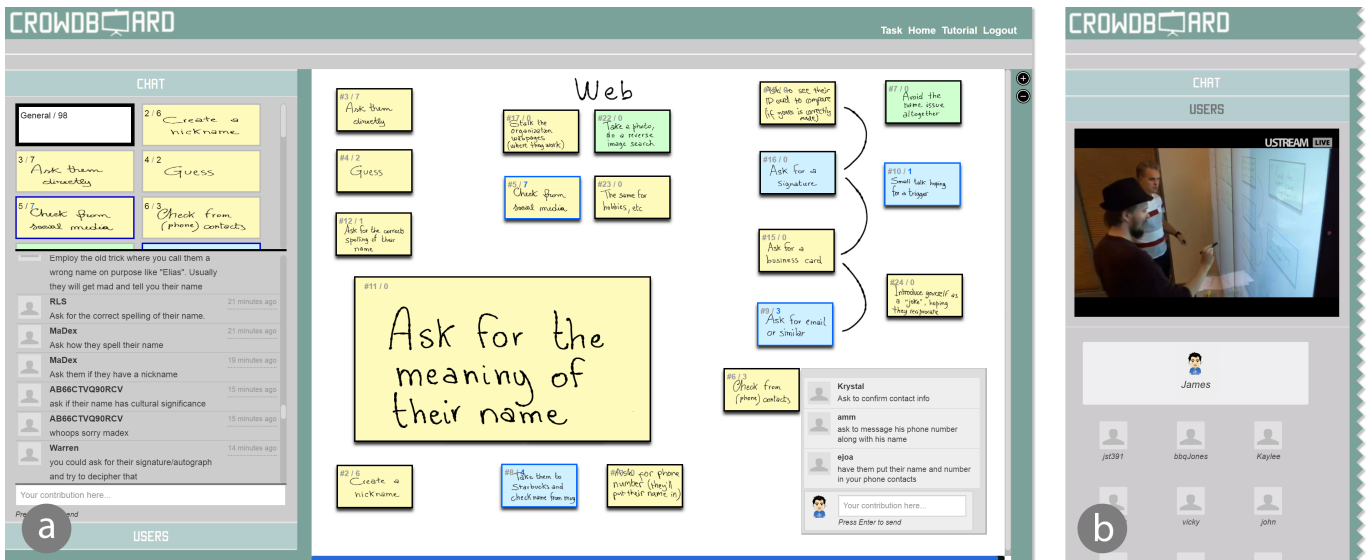


Figure 7. Crowdboard 2.0 (web interface). a) The main interface. Crowd workers see a list of the design conversations (upper left) and whiteboard activity (right). They can leave comments, which, in turn, are shown on the digital board. b) Users tab. This view shows the live stream of the meeting and the list of crowd workers involved in the discussion.

left corner of a note denote the note's order of creation and the number of comments attached to it. Clicking on those numbers expands or collapses the chat attached to the note. In contrast to Crowdboard 1.0, where crowd workers created discussion threads to refer to specific ideas on the board, in version 2.0, the threads are attached to specific ideas from the moment that the in-person ideators write down those ideas. When ideas are moved on the board, the corresponding comments follow the ideas automatically, unlike in Crowdboard 1.0. This design choice, while helping maintaining consistency, also leads to better organization of ideas and eliminates the problem of duplicate discussion threads that refer to the same ideas.

The left panel (Figure 5a) contains the general chat, where crowd ideators can add comments and ideas that do not relate to anything present on the board. From the left panel, it is also possible to access chats related to specific notes by selecting those notes from the overview at the top of the panel. The panel content automatically switches back to show the general chat after a few seconds of inactivity.

Crowdboard 2.0 includes several improvements to the notifications. When a new comment is added to a note, it is immediately shown on top of that note for a few seconds before fading away. When the new comment is added, the borders of the corresponding note blink in a blue color. The color and thickness of a note's borders indicate its status: Black is the default, blue indicates that it contains unread comments, and thick black means that the note is currently selected. Unread comments, such as those in a long chat that the user needs to scroll to the bottom of, have small badge that says "new" attached. When the comments become visible, the badge fades away after a few seconds.

The button in the top-right corner of the interface gives access to a menu containing the session- and crowd-management widgets (Figure 5e). The crowd-recruitment widget includes a few

options, such as specifying the time at which crowds should join a live discussion, the time allotted for the discussion, the number of workers to be recruited, and the money that will be paid to each worker. The crowd-review widget permits users to go through a summary of the contributions that the provided during the discussion and to easily identify possibly problematic crowd workers. The same review widget is used to confirm the payment for the contributing crowd workers.

Web interface

The Crowdboard 2.0 web interface (Figure 7) reflects the new features introduced in the studio interface. As in the previous version of the system crowd ideators can hear and see the ongoing brainstorming session, which is led by in-person ideators working with the studio interface. The crowd ideators see a synchronously updated view of the workspace, including sketches and virtual notes, and they can contribute by adding comments and ideas to specific notes or to the general chat. The workspace content that is sent to the crowd is scaled to fit various screen sizes. Because the screens of the crowd workers' computers are likely to be smaller than the big screen used in the studio, we added zooming features to the interface. A note can be expanded or collapsed by clicking on it, and the user can zoom in and out the entire workspace using the "+" and "-" buttons on the top right corner of the main panel. Notes in the web interface have the same behavior as those in the studio interface, but they cannot be edited.

STUDY 2: CROWDBOARD'S EFFECT ON IDEATION

We evaluate Crowdboard 2.0 using a quantitative controlled experiment with in-person teams of two people generating ideas either using real-time support from online crowd workers or on their own, without crowd involvement. We hypothesize that Crowdboard will improve teams' creativity by enabling them to sustain higher levels of creative output over time, compared to unaided teams can.

Method

Participants

This study leverages two populations of participants to serve two roles in the study: 1) in-person teams of ideators and 2) crowd ideators. We recruited 40 people (12 females) from the computer science departments of two universities to work in teams of two. The average age of these participants was 27.60 [SD = 5.48]. Each received one movie ticket as compensation for participation. We also recruited 143 workers (40.4% females, mean age = 33.88 [SD = 9.07]) from Mechanical Turk to participate as crowd ideators. The crowd ideators teamed up with the in-persons teams in one of the experimental conditions. Overall, 10 sessions were conducted for the crowd-aided condition yielding an average of 14.3 crowd ideators per session. We limited recruitment to workers who had completed at least 1,000 HITs and who had an approval rate greater than 95%. Crowd workers were paid \$3.00 for their participation.

Study Design

The evaluation was conducted as a between-subjects experiment. In each trial, dyads of in-person ideators were randomly assigned to one of two conditions:

1. In the **Crowdboard** condition, in-person ideators generated ideas using the Crowdboard 2.0 studio interface with crowd support;
2. In the **whiteboard only** condition, ideators generated ideas using the Crowdboard 2.0 studio interface, without crowd support;

Brainstorming Prompt

Participants generated ideas for the following problem: “*Imagine you are in a social setting and you’ve forgotten the name of someone you know. How might you recall their name without directly asking them? Be as specific as possible in your descriptions.*” This problem has already been validated in other studies of crowd-based idea generation [11].

Procedure

In-person ideators went through the following procedure. After obtaining the participants’ informed consent, the experimenter used the Crowdboard 2.0 studio interface to invite 20 crowd workers to join the brainstorming session, which was to start in 20 minutes from that moment. The experimenter then explained the overall task and handed out instructions for brainstorming and a modified version of Osborn’s [30] brainstorming rules: 1) criticism is ruled out; 2) freewheeling is welcomed; 3) quantity is wanted; 4) combination and improvements are sought; and 5) users stay focused on the task. Drawing on the lessons learned from Study 1 we also provided the instruction to remember to acknowledge the crowd.

In-person ideators then went through a training session with simulated crowd input. After this, they performed a warm-up brainstorming session on “alternative uses for a pencil.” Then, the experimenters showed the participants the brainstorming topic and allowed for a couple of minutes of individual reflection. After that, the experimenters used the “Go live” function of Crowdboard 2.0 to start streaming the conversation to the crowd. When the session was over, the experimenters stopped

the live stream and asked the in-person participants to fill out a survey. After completing the survey, the experimenters asked the participants to read the list of crowd contributions and mark which ones represented valid ideas (rather than just comments). For the contributions marked as valid ideas, the participants also marked whether 1) they saw it; 2) they used it; and 3) it was (or would have been) somewhat useful to see it.

Crowd ideators went through a different procedure. Once launched into Crowdboard from Mechanical Turk, the crowd ideators provided informed consent. The ideators then watched a video tutorial, including instructions about the roles of both general chat and sticky note-related chats (see Figure 5a and Figure 5d). After completing the tutorial, the ideators waited for the brainstorming session to begin at the scheduled time. In the meantime, they were allowed to perform other tasks. At the scheduled time, the crowd ideators were automatically presented with the brainstorming topic. At the same moment they started hearing and seeing the in-person ideators working on the studio interface. After 20 minutes the session was over and the crowd ideators were automatically directed to a short survey with questions about demographic information and their experiences during the brainstorming session.

Dependent Measures

Novelty

Novelty was conceptualized as the degree to which an idea surprised a judge. We recruited 15 workers from Amazon MTurk to rate the ideas. Each worker rated a set of 70 ideas, and 3 raters judged each idea. Inter-rater reliability was substantial, ICC(2,3) = 0.68.

Value

Value was conceptualized as the estimated likelihood that the idea would work if that it were actually implemented. Each worker who rated novelty also rated value for the same set of ideas. Inter-rater reliability was acceptable, ICC(2,3) = 0.57.

Creativity

Following prior theoretical and empirical work [11, 34], we operationalized *creativity* as the product of novelty and value. This captures the intuition that creative ideas should be high on *both* novelty and value.

Results

The in-person ideators across both conditions generated 325 ideas (with 22 invalid ideas removed), 173 in the Crowdboard condition ($M = 17.3$ per session, $SD = 7.36$), and 152 in the control condition ($M = 15.2$ per session, $SD = 3.27$). In the Crowdboard condition, crowd ideators sent 783 contributions ($M = 78.3$ per session, $SD = 24.1$), 530 of which ($M = 53.0$ per session, $SD = 19.1$) the in-person ideators rated as valid; this corresponds to 68% of all contributions. Table 1 shows descriptive statistics regarding the Crowdboard condition (averaged across sessions). Out of an average of 53 valid ideas, 29.7 were attached to sticky note chats (SNC) and 23.3 were posted to the general chat (GC). According to the post-study questionnaire results, in-person ideators noticed 74% of all

Table 1. Crowd-generated contributions. SNC = Sticky Note Chats; GC = General Chat

	SNC		GC		Total	
	M	SD	M	SD	M	SD
Contributions	45.5	16.2	32.8	16.6	78.3	24.1
Ideas	29.7	9.1	23.3	15.6	53.0	19.1
Seen	25.8 (87%)	8.5	13.6 (58%)	9.6	39.4 (74%)	12.8
Useful	22.9 (77%)	10.7	16.9 (73%)	9.8	39.8 (75%)	15.1
Used	9.6 (32%)	6.6	8.2 (35%)	5.8	17.8 (34%)	9.8
Useful not seen	2.5 (11%)	3.6	5.8 (34%)	6.6	8.3 (21%)	9.6

ideas posted on the board, including 87% of SNC ideas and 58% of GC ideas. Additionally 75% of all ideas were rated as useful. Of those, 21% were not noticed (11% of SNC and 34% of GC).

To investigate the impact that Crowdbord had on the groups' ideations, we estimate a series of linear mixed effects models for each dependent measure, with random effects of session (to account for nonindependence of ideas within sessions). The idea is the unit of analysis. Our baseline models only have a fixed effect of time in addition to the random effects, to capture the known effect that ideas tend to get better over time during brainstorming [5]. We then test Crowdbord by adding a fixed effect of condition (Crowdbord vs. baseline) to the model. We test for significance of the Crowdbord intervention in two ways. First, since recent work [1] has shown that testing for significance of a single fixed effect in a model is only reliable for very large datasets, we use a Likelihood ratio test, where we compare a more complex model (with the fixed effect added) to a baseline model (without the fixed effect) and test determine whether there is a statistically significant improvement in model fit to the data (which is measured by the degree to which the model deviates from the data; lower is better). Second, to guard against overfitting, we test whether the improved fit comes at the cost of unnecessarily increased model complexity (which might indicate overfitting, and consequently, poor generalizability of the model); to accomplish this, we compare models in terms of their Akaike Information Criterion (AIC [2]), which adds a penalty for the number of parameters (essentially penalizing complex models) to the estimated model fit. All models are estimated in R, using the `glmer` function in the `lme4` package.

Novelty

Model #1 has *Time* as the only fixed effect. This model yields a statistically significant gain in model fit over the null model, Likelihood ratio $\chi^2(1) = 24.54$, $p < .001$, and a reduction in Akaike Information Criterion (AIC) (from 1104.1 to 1081.6), indicating that this model is a better fit to the data (without overfitting). In this model, *Time* is a significant positive predictor of novelty, such that each additional minute yields an expected increase in idea novelty of 0.06 points, 95% $CI = [0.04, 0.08]$, $p < .001$. In other words, ideas get more novel over time (consistent with the serial order effect in prior literature [5]).

In model #2, we add the hypothesized *condition* \times *time* interaction as a fixed effect. This model does not yield a statistically significant gain in model fit over model #1, Likelihood ratio $\chi^2(1) = 0.35$, $p = .56$, and an increase in AIC (from 1081.6 to 1083.2), indicating that this model is a poorer fit to the data than model #1. Thus, we conclude that novelty increases over time, but Crowdbord does not influence this trajectory.

Value

Model #1 has *Time* as the only fixed effect. This model yields a statistically significant gain in model fit over the null model, Likelihood ratio $\chi^2(1) = 12.48$, $p < .001$, and a reduction in AIC (from 1087.0 to 1076.5), indicating that this model is a better fit to the data (without overfitting). In this model, *Time* is a significant negative predictor of value, such that each additional minute yields an expected decrease in idea value of 0.04 points, 95% $CI = [-0.07, -0.02]$, $p < .001$. In other words, ideas get less valuable over time.

In model #2, we add the hypothesized *condition* \times *time* interaction as a fixed effect. This model does not yield a statistically significant gain in model fit over model #1, Likelihood ratio $\chi^2(1) = 2.17$, $p = .14$, although it does yield a slight decrease in AIC (from 1076.5 to 1076.3), suggesting that this model may provide useful information over and above model #1. In this model, the interaction parameter for Crowdbord is positive, suggesting that Crowdbord use might slightly slow the decline of value over time, $B = 0.02 [0.01, 0.05]$, but the coefficient for the interaction is not statistically significant, $p = .16$. This result should therefore be interpreted with caution.

Creativity

Table 2 shows the results of the models for creativity. Model #1 has *Time* as the only fixed effect. This model does not yield a statistically significant gain in model fit over the null model, Likelihood ratio $\chi^2(1) = 1.67$, $p = .20$, and an increase in AIC (from 1999.1 to 1999.4), indicating that this model is a better fit to the data (without overfitting). This suggests that, overall, there is not a simple relationship between time and

Table 2. Summary of linear mixed effects models for creativity of ideas. There is no significant main effect of time in session on creativity of ideas, but a significant interaction with experimental condition, such that creativity only increases over time for Crowdbord ideas.

	#1: Time	#2: Add Crowdbord interaction
Fixed effects		
Intercept	13.21 [11.85, 14.58]	13.23 [11.91, 14.56]
Time	0.06 [-0.03, -0.16] **	-0.02 [-0.15, 0.10] **
Crowdbord \times time		0.16 [0.02, 0.30] *
Random effects		
Variance across sessions for intercept	4.07	3.58
Residual	24.98	24.80
Model fit statistics		
Deviance	1099.40	1986.40
AIC	1099.40	1996.40

* $p < .05$; ** $p < .01$; *** $p < .001$; 95% CI (Wald) = [lower, upper].

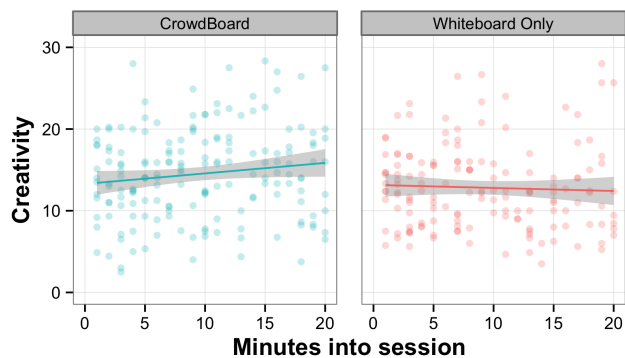


Figure 8. Study 2. Ideas increase in creativity over time, at a higher rate with the Crowdboard condition.

creativity. This is not surprising given the opposing effects of time on novelty and value, and the negative correlation between novelty and value. It stands to reason that, with no other information, we should expect to see no change in creativity over time (as novelty increases but value decreases).

However, model #2, which adds the hypothesized *condition* \times *time* interaction as a fixed effect, does yield a significant gain in model fit over the null model, Likelihood ratio $\chi^2(2) = 6.71$, $p < .05$, and a reduction in AIC (from 1999.1 to 1996.4). This model still has no main effect of time, but does have a statistically significant interaction between Crowdboard and time, $B = 0.17$ [0.02, 0.30], $p < .05$. Inspection of Figure 8 provides a qualitative interpretation of this interaction: under normal conditions (Whiteboard only), ideas do not get more creative over time; however, with Crowdboard, ideas *do* get more creative over time.

Discussion

This research provides evidence that real-time crowds can positively affect idea generation sessions by providing a continuous flow of creative ideas that fuels the conversation. With Crowdboard ideas *do* get more creative over time, although the effect is modest in size. One possible explanation is that results from the mutual stimulation of associations caused by the real-time interactions between in-person and crowd ideators. In contrast to other asynchronous approaches, here, people can build on each other’s ideas in real-time and engage in stimulating discussions with other participants, all while avoiding the repetition of obvious and mundane ideas.

Though quality control has traditionally been a concern for this kind of high-level creative contribution from a crowd, Crowdboard experiments have shown that senior Mechanical Turk workers may be a reliable source of inspiration. In our studies we saw no clear instance of offensive or otherwise problematic input that required moderation. The reputation mechanism that senior workers use is itself an incentive for good behaviour [19]. Additionally the Mechanical Turk participants considered the task to be engaging and enjoyable—a factor that could have contributed to making them less prone to opportunistic behaviors.

Our experiments point out that the design choice of having SNC could provide some advantages over more traditional chat styles, such as GC. In Study 2, the in-person ideators missed most of the valid ideas that were added to the general chat (only 58% were noticed), but they noticed a good percentage of those attached to sticky notes (87%). Another possible advantage of SNC would be the organization of ideas into logical units. However in Study 2 SNC were not always used as intended. We observed that ideas were not necessarily posted to the most relevant SNC; instead they were often posted to the currently active one or to the most popular one, with the users obviously intending to maximize the chances that the idea would be seen in the shorter amount of time possible. Although this behavior is not necessarily counterproductive—and might have contributed to the better percentage of the ideas being noticed—it may be reasonable to assume that the intended usage of SNC may have been even more advantageous. Future versions of the system should consider the inclusion of measures to address the problem.

Crowdboard is meant to be used as a stand-alone tool that does not require a big organizational effort. For example, it can easily handle crowd recruitment and payment, thanks to the integrated crowd management widgets. However Crowdboard sessions may still require the involvement of a moderator. In our studies the system showed the potential for self-moderation, as most concerns and doubts raised in the general chat were promptly addressed by other crowd participants. Nevertheless moderator’s intervention was sometimes needed when off-topic discussions on individual technical issues were causing too much distraction from the main task. Future iterations of the system may include different crowd roles, including moderators, and/or features for self-moderation [14].

CONCLUSIONS

Synchronous interaction among in-person and crowd ideators offers the opportunity to boost early stage design processes without the organizational complexity and time cost of large-scale innovation strategies. This paper presents Crowdboard, a novel whiteboard system that allows a team of in-person ideators to collaborate in real-time with online crowds on creative problems. With two user studies we tested the feasibility of the idea and the effects of real-time crowds on ideation. Results showed that Crowdboard helped participants generate ideas that increased in creativity over time at a higher rate than they did when using the whiteboard without a crowd. While there are still many open challenges regarding real-time interactions among in-person and crowd ideators, our work suggests that this approach could be a fast and effective way to boost early-stage ideation, and provides pointers for future developments of the paradigm.

ACKNOWLEDGMENTS

This research was partially funded by TEKES (Re:Know2) and the National Science Foundation (awards 1122206 and 1122320). We thank Daniel Lee, Peeyush Goyal, Antonio Gentile, and Haakon Faste for their valuable input during the earliest stages of the project, and Walter Lasecki for his valuable feedback on early drafts of the paper.

REFERENCES

1. Alan Agresti. 2007. *An Introduction to Categorical Data Analysis* (2nd ed.). Wiley-Interscience.
2. Hirotugu Akaike. 1998. *Information Theory and an Extension of the Maximum Likelihood Principle*. Springer New York, New York, NY, 199–213. DOI: http://dx.doi.org/10.1007/978-1-4612-1694-0_15
3. Salvatore Andolina, Khalil Klouche, Diogo Cabral, Tuukka Ruotsalo, and Giulio Jacucci. 2015. InspirationWall: Supporting Idea Generation Through Automatic Information Exploration. In *Proceedings of the 2015 ACM SIGCHI Conference on Creativity and Cognition (C&C '15)*. ACM, New York, NY, USA, 103–106. DOI: <http://dx.doi.org/10.1145/2757226.2757252>
4. Brian P. Bailey and Eric Horvitz. 2010. What's Your Idea?: A Case Study of a Grassroots Innovation Pipeline Within a Large Software Company. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*. ACM, New York, NY, USA, 2065–2074. DOI: <http://dx.doi.org/10.1145/1753326.1753641>
5. Roger E. Beaty and Paul J. Silvia. 2012. Why do ideas get more creative across time? An executive interpretation of the serial order effect in divergent thinking tasks. *Psychology of Aesthetics, Creativity, and the Arts* 6, 4 (2012), 309–319. DOI: <http://dx.doi.org/10.1037/a0029171>
6. Michael S. Bernstein, Joel Brandt, Robert C. Miller, and David R. Karger. 2011. Crowds in Two Seconds: Enabling Realtime Crowd-powered Interfaces. In *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology (UIST '11)*. ACM, New York, NY, USA, 33–42. DOI: <http://dx.doi.org/10.1145/2047196.2047201>
7. Michael S. Bernstein, Greg Little, Robert C. Miller, Björn Hartmann, Mark S. Ackerman, David R. Karger, David Crowell, and Katrina Panovich. 2015. Soylent: A Word Processor with a Crowd Inside. *Commun. ACM* 58, 8 (July 2015), 85–94. DOI: <http://dx.doi.org/10.1145/2791285>
8. Jeffrey P. Bigham and Walter S. Lasecki. 2014. Crowd Storage: Storing Information on Existing Memories. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 601–604. DOI: <http://dx.doi.org/10.1145/2556288.2557159>
9. Kevin J Boudreau and Karim R Lakhani. 2013. Using the crowd as an innovation partner. *Harvard business review* 91, 4 (2013), 60–69.
10. Joel Chan, Steven Dang, and Steven P. Dow. 2016a. Comparing Different Sensemaking Approaches for Large-Scale Ideation. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 2717–2728. DOI: <http://dx.doi.org/10.1145/2858036.2858178>
11. Joel Chan, Steven Dang, and Steven P. Dow. 2016b. Improving Crowd Innovation with Expert Facilitation. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW '16)*. ACM, New York, NY, USA, 1223–1235. DOI: <http://dx.doi.org/10.1145/2818048.2820023>
12. Michael Diehl and Wolfgang Stroebe. 1987. Productivity loss in brainstorming groups: Toward the solution of a riddle. *Journal of Personality and Social Psychology* (1987), 497–509. DOI: <http://dx.doi.org/10.1037/0022-3514.53.3.497>
13. Steven Dow, Elizabeth Gerber, and Audris Wong. 2013. A Pilot Study of Using Crowds in the Classroom. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, USA, 227–236. DOI: <http://dx.doi.org/10.1145/2470654.2470686>
14. Arpita Ghosh, Satyen Kale, and Preston McAfee. 2011. Who Moderates the Moderators?: Crowdsourcing Abuse Detection in User-generated Content. In *Proceedings of the 12th ACM Conference on Electronic Commerce (EC '11)*. ACM, New York, NY, USA, 167–176. DOI: <http://dx.doi.org/10.1145/1993574.1993599>
15. François Guimbretière, Maureen Stone, and Terry Winograd. 2001. Fluid Interaction with High-resolution Wall-size Displays. In *Proceedings of the 14th Annual ACM Symposium on User Interface Software and Technology (UIST '01)*. ACM, New York, NY, USA, 21–30. DOI: <http://dx.doi.org/10.1145/502348.502353>
16. Raja Gumienny, Lutz Gericke, Matthias Wenzel, and Christoph Meinel. 2013. Supporting creative collaboration in globally distributed companies. In *Proc. of the 2013 conference on Computer supported cooperative work (CSCW '13)*. ACM, 995–1007. DOI: <http://dx.doi.org/10.1145/2441776.2441890>
17. Joshua Hailpern, Erik Hinterbichler, Caryn Leppert, Damon Cook, and Brian P. Bailey. 2007. TEAM STORM: Demonstrating an Interaction Model for Working with Multiple Ideas During Creative Group Work. In *Proceedings of the 6th ACM SIGCHI Conference on Creativity & Cognition (C&C '07)*. ACM, New York, NY, USA, 193–202. DOI: <http://dx.doi.org/10.1145/1254960.1254987>
18. Otmar Hilliges, Lucia Terrenghi, Sebastian Boring, David Kim, Hendrik Richter, and Andreas Butz. 2007. Designing for Collaborative Creative Problem Solving. In *Proceedings of the 6th ACM SIGCHI Conference on Creativity & Cognition (C&C '07)*. ACM, New York, NY, USA, 137–146. DOI: <http://dx.doi.org/10.1145/1254960.1254980>
19. Audun Jsang, Roslan Ismail, and Colin Boyd. 2007. A survey of trust and reputation systems for online service provision. *Decision Support Systems* 43, 2 (2007), 618 – 644. DOI: <http://dx.doi.org/10.1016/j.dss.2005.05.019>

20. Martin Kaltenbrunner, Till Bovermann, Ross Bencina, and Enrico Costanza. 2005. TUIO: A protocol for table-top tangible user interfaces. In *Proc. of the The 6th Int'l Workshop on Gesture in Human-Computer Interaction and Simulation*. 1–5.
21. Joy Kim, Justin Cheng, and Michael S. Bernstein. 2014. Ensemble: Exploring Complementary Strengths of Leaders and Crowds in Creative Collaboration. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW '14)*. ACM, New York, NY, USA, 745–755. DOI: <http://dx.doi.org/10.1145/2531602.2531638>
22. Aniket Kittur. 2010. Crowdsourcing, Collaboration and Creativity. *XRDS* 17, 2 (Dec. 2010), 22–26. DOI: <http://dx.doi.org/10.1145/1869086.1869096>
23. Scott R. Klemmer, Mark W. Newman, Ryan Farrell, Mark Bilezikjian, and James A. Landay. 2001. The Designers' Outpost: A Tangible Interface for Collaborative Web Site. In *Proceedings of the 14th Annual ACM Symposium on User Interface Software and Technology (UIST '01)*. ACM, New York, NY, USA, 1–10. DOI: <http://dx.doi.org/10.1145/502348.502350>
24. Aaron Michael Koblin. 2010. The Johnny Cash Project. (2010). <http://www.thejohnnycashproject.com/>.
25. Walter S. Lasecki, Juho Kim, Nick Rafter, Onkur Sen, Jeffrey P. Bigham, and Michael S. Bernstein. 2015. Apparition: Crowdsourced User Interfaces That Come to Life As You Sketch Them. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 1925–1934. DOI: <http://dx.doi.org/10.1145/2702123.2702565>
26. William J. Baumol Lu Hong, Scott E. Page. 2004. Groups of Diverse Problem Solvers Can Outperform Groups of High-Ability Problem Solvers. *Proceedings of the National Academy of Sciences of the United States of America* 101, 46 (2004), 16385–16389. <http://www.jstor.org/stable/3373827>
27. Kurt Luther, Casey Fiesler, and Amy Bruckman. 2013. Redistributing Leadership in Online Creative Collaboration. In *Proceedings of the 2013 Conference on Computer Supported Cooperative Work (CSCW '13)*. ACM, New York, NY, USA, 1007–1022. DOI: <http://dx.doi.org/10.1145/2441776.2441891>
28. Kurt Luther, Jari-Lee Tolentino, Wei Wu, Amy Pavel, Brian P. Bailey, Maneesh Agrawala, Björn Hartmann, and Steven P. Dow. 2015. Structuring, Aggregating, and Evaluating Crowdsourced Design Critique. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW '15)*. ACM, New York, NY, USA, 473–485. DOI: <http://dx.doi.org/10.1145/2675133.2675283>
29. MultiTaction. 2017. MultiTaction — MT Cell. (2017). Retrieved March 7, 2017 from <https://www.multitaction.com/hardware/mt-cell>.
30. Alex F Osborn. 1953. Applied imagination. (1953).
31. Cheong Ha Park, KyoungHee Son, Joon Hyub Lee, and Seok-Hyung Bae. 2013. Crowd vs. Crowd: Large-scale Cooperative Design Through Open Team Competition. In *Proceedings of the 2013 Conference on Computer Supported Cooperative Work (CSCW '13)*. ACM, New York, NY, USA, 1275–1284. DOI: <http://dx.doi.org/10.1145/2441776.2441920>
32. Paul B Paulus and Huei-Chuan Yang. 2000. Idea Generation in Groups: A Basis for Creativity in Organizations. *Organizational Behavior and Human Decision Processes* 82, 1 (2000), 76 – 87. DOI: <http://dx.doi.org/10.1006/obhd.2000.2888>
33. Pao Siangliulue, Joel Chan, Steven P. Dow, and Krzysztof Z. Gajos. 2016. IdeaHound: Improving Large-scale Collaborative Ideation with Crowd-Powered Real-time Semantic Modeling. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology (UIST '16)*. ACM, New York, NY, USA, 609–624. DOI: <http://dx.doi.org/10.1145/2984511.2984578>
34. Dean Keith Simonton. 2016. Defining Creativity: Don't We Also Need to Define What Is Not Creative? *The Journal of Creative Behavior* (Jan. 2016), n/a–n/a. DOI: <http://dx.doi.org/10.1002/jocb.137>
35. Jasjit Singh and Lee Fleming. 2010. Lone Inventors as Sources of Breakthroughs: Myth or Reality? *Management Science* 56, 1 (2010), 41–56. DOI: <http://dx.doi.org/10.1287/mnsc.1090.1072>
36. James Surowiecki. 2005. *The wisdom of crowds*. Anchor.
37. Eliot Van Buskirk. 2010. Google Struggles to Give Away \$10 Million. (2010). <http://www.wired.com/business/2010/06/google-struggles-to-give-away-10-million/all/>
38. Hao-Chuan Wang, Dan Cosley, and Susan R. Fussell. 2010. Idea Expander: Supporting Group Brainstorming with Conversationally Triggered Visual Thinking Stimuli. In *Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Work (CSCW '10)*. ACM, New York, NY, USA, 103–106. DOI: <http://dx.doi.org/10.1145/1718918.1718938>
39. Andy Warr and Eamonn O'Neill. 2005. Understanding Design As a Social Creative Process. In *Proceedings of the 5th Conference on Creativity & Cognition (C&C '05)*. ACM, New York, NY, USA, 118–127. DOI: <http://dx.doi.org/10.1145/1056224.1056242>
40. Anbang Xu, Shih-Wen Huang, and Brian Bailey. 2014. Voyant: Generating Structured Feedback on Visual Designs Using a Crowd of Non-experts. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW '14)*. ACM, New York, NY, USA, 1433–1444. DOI: <http://dx.doi.org/10.1145/2531602.2531604>

41. Lixiu Yu, Aniket Kittur, and Robert E. Kraut. 2014. Distributed Analogical Idea Generation: Inventing with Crowds. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 1245–1254. DOI: <http://dx.doi.org/10.1145/2556288.2557371>
42. Lixiu Yu, Aniket Kittur, and Robert E. Kraut. 2016. Encouraging “Outside- The- Box” Thinking in Crowd Innovation Through Identifying Domains of Expertise. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW '16)*. ACM, New York, NY, USA, 1214–1222. DOI: <http://dx.doi.org/10.1145/2818048.2820025>
43. M. P. Zanna, J. M. Olson, W. Stroebe, B. A. Nijstad, and E. F. Rietzschel. 2010. Beyond Productivity Loss in Brainstorming Groups: The Evolution of a Question. In *Advances in Experimental Social Psychology*. Vol. 43. San Diego, CA, 157–203.