

The Effect of Social Interaction on Facilitating Audience Participation in a Live Music Performance

Sang Won Lee
Computer Science
Virginia Tech
sangwonlee@vt.edu

Aaron Willette
Performing Arts Technology
University of Michigan
aawill@umich.edu

Danai Koutra
Computer Science & Engineering
University of Michigan
dkoutra@umich.edu

Walter S. Lasecki
Computer Science & Engineering
University of Michigan
wlasecki@umich.edu

ABSTRACT

Facilitating audience participation in a music performance brings with it challenges in involving non-expert users in large-scale collaboration. A musical piece needs to be created live, over a short period of time, with limited communication channels. To address this challenge, we propose to incorporate social interaction through mobile music instruments that the audience is given to play with, and examine how this feature sustains and affects the audience involvement. We test this idea with an audience participation music system, *Crowd in C*. We realized a participation-based musical performance with the system and validated our approach by analyzing the interaction traces of the audience at a performance. The result indicates that the audience members were actively engaged throughout the performance, with multiple layers of social interaction available in the system. We also present how the social interactivity among the audience shaped their interaction in the music making process.

Author Keywords

Audience Participation, Mobile Music, Music Performances

CCS Concepts

•Applied computing → Sound and music computing; Performing arts; •Human-centered computing → Collaborative content creation;

INTRODUCTION

Audience participation in a public event is an effective way to engage large-scale audiences and create an artifact collaboratively [28, 13, 19, 42]. It is, however, challenging due to its instantaneous nature and the effort needed to coordinate

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 ACM 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 '19, June 23–26, 2019, San Diego, CA, USA

© 2019 ACM. ISBN 978-1-4503-5917-7/19/06...\$15.00

DOI: <https://doi.org/10.1145/3325480.3325509>

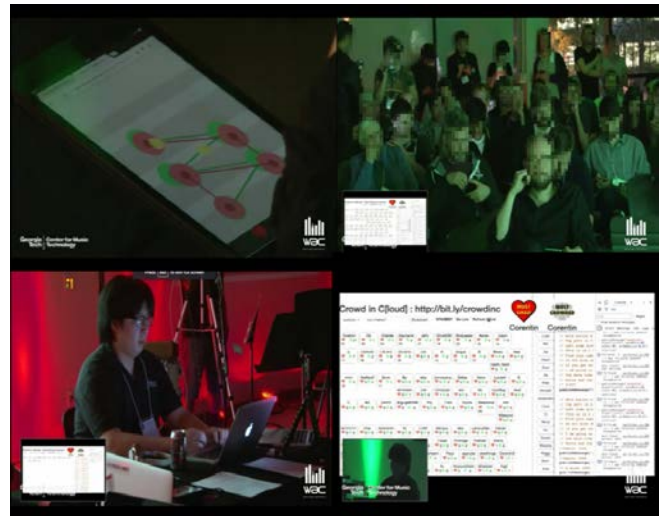


Figure 1. *Crowd in C* in Action. *Crowd in C* has been performed five times, involving a total of more than 350 audience members. Social interaction via musical instruments is embedded in the instrument to facilitate audience engagement and the creative process.

audiences on that scale. In particular, researchers have explored audience participation performances in live concerts using audiences' smartphones as musical instruments [33, 35]. Audience participation in a concert comes with a challenge of sustaining participants' interest to continue participating, especially given that the outcome of the event relies heavily on the audience's active participation. In addition, the duration of the designed participatory experience tends to be much shorter than other types of events (tens of minutes at most per piece, usually) and only non-verbal communication is allowed. Therefore, there exists a multitude of challenges for musicians in organizing audience participation in a concert: 1) lowering the barriers of involving a common audience in a highly interactive artistic practice, 2) engaging the audience in a participatory, collaborative creation, and 3) safeguarding the quality of the artifact (i.e., the musical piece). We explored these challenges in a recent audience participation

performance with an interactive system, *Crowd in C* [33]. (Figure 1.) The system supports a hybrid event [41] in which audience members create an artifact (e.g., the sound of music) in a co-located setup while digitally connected over a social network provided by the system. In this paper, we particularly focus on the second challenge—engaging the audience in the participatory music-making.

We hypothesize that promoting social interactivity through computer-mediated communication channels in an interactive music system can effectively sustain audience engagement and help them be part of a large-scale participatory music performance. In addition, we are interested in understanding the effect of social interactivity on creative engagement with participatory music performance. Creative engagement is a concept for interactive art, which raises challenges like how to anticipate and define the user experience, and it has been mostly explored in the context of exhibitions/galleries [14]. We examine the effect of social interaction through an interactive music piece called, *Crowd in C*, in which audience members can interact with each other as in social media (e.g., browsing, endorsing, real-time interaction). Specific research questions that we seek to answer are as follows:

- **RQ1—Sustaining Audience Engagement:** Does incorporating social interaction among audience members in an interactive system sustain the participatory activity in a large-scale, participatory musical performance?
- **RQ2—Understanding the effect of social interactivity in the audience’s creative engagement:** How does social interactivity influence and shape the ways in which the audience interacts with the interactive music system and participates in an interactive music performance?

We validate our approach by analyzing a recent performance of *Crowd in C*. We begin with a discussion of the background and related work in facilitating audience participation events in music and other domains. We then discuss how *Crowd in C* works, and we demonstrate how social interaction within the system is designed to help audience members participate in music-making. We then present the result of our data analysis of interaction traces from a recent performance to validate our approach. Lastly, we conclude the paper with further discussion on the study and future works. The findings from this study will contribute to the existing knowledge of audience experience research in informing artists and practitioners, and advancing our understanding in creative engagement particularly for interactive music performance [14, 3, 55].

RELATED WORK

This work builds on previous audience participation work in which audiences are engaged in a public event, in various settings. In the rest of this section, we introduce both general events and those in a musical setting.

Audience Participation in Hybrid Events

Live participation has emerged as a new topic in HCI, in which researchers and practitioners utilize computer-mediated communication to produce engaging public events in a co-located

setup [42]. Such tools can be used in various settings—large-scale classrooms [2, 12], seminar-style group discussions [18], panel discussions [19], creating large-scale structures [28], theater performances [5], collective games [4], and art galleries [22]. In general, there are three methods of facilitating audience participation. One frequently used method is to collect votes using various apparatus—so-called audience response systems—including, but not limited to, clickers [24], visual markers [7], glowing sticks [17], paddles [4], and audience-owned devices [50]. Another approach is to enable a backchannel in which audience members can communicate with the event organizers (or *performers*) and other audience members [12, 19, 2, 56]. Another approach is to allow a limited number of audience members to individually participate in an open space directly, one at a time, and over time, rather than enabling large-scale concurrent actions [22, 46]. Our approach is similar to enabling backchannels in which audience members can interact with each other in the channel provided by musical instruments.

Another opportunity afforded by audience participation has a rich potential for fostering creativity by involving the audience directly in the process of artifact creation in a public event. Supporting the creation of an (often physical) artifact in a physical location using technology presents a set of challenges unique to its hybrid (digital + physical) setting [41, 6, 28]. This is because the activity of artifact creation, as opposed to other activities like open discussion, needs to be coordinated and guided precisely with a set of technologies [16, 28]. Therefore, we need to better understand how to design a participatory platform on which artists can create and deliver artistic artifacts with the audience. Live music exemplifies such challenges, as music-making is a highly interactive practice which requires a lot of pre-production effort, and the physical (audible) artifact is a temporal work of art delivered *during its creation*.

Audience Participation in Musical Performances

Musicians and researchers have attempted to engage audiences by directly involving them in the music-making process at live performances. For example, *We Will Rock You* by British rock band Queen is famous for the repetition of a simple participatory rhythm—*stomp, stomp, clap*—which results in successful participation. It has been used frequently at sports events by fans supporting their teams.

One of the biggest challenges of enabling audience participation in live concerts concerns how one can reconcile the musicians’ desire to deliver a compelling piece of music while the audience directly influences the music. For example, if one decides to distribute traditional acoustic instruments (e.g., violins) and asks an audience to participate in performing a piece they have never heard of, it is likely that the audience would hesitate to participate in the performance; even if the audience does not hesitate, the sound is likely to interfere with the artifact that the organizing musicians wish to produce [32]. To that end, musicians develop devices that can facilitate audience participation from multiple angles—composition, pre-production, apparatus, and guiding gestures in live performance. In the example of *We Will Rock You*, the composers designed the participatory actions to be repetitive

and simple for non-experts to perform. Through gestures and drum patterns, the band's efforts coordinate the audience's participation. In a contemporary musical piece, *Moths* by Jean Hasse, a conductor used a graphical, open-form score. By guiding people to read a drawing and interpret it as pitch over time, the audience is led to whistle in accordance with the conductor's gestures [20]. Another example of this coordination can be found in Bobby McFerrin's improvisation. During his talk at the World Science Festival 2009, McFerrin demonstrated audience participation, guiding his audience to sing the right note as he hopped around the stage¹ [40]. Some early examples included state-of-the-art technologies of their time (such as cassette tapes from *The Flaming Lips'* The Parking Lot Experiments [11], or Maseda's radios from [39]).

Computer-mediated audience participation can address the challenges inherent to participatory music performances. One approach is to allow the audience to take on the composer's role, indirectly influencing the music without directly producing any sound. *Glimmer* exemplifies this approach well: audience members wave light sticks to influence the music, and the movement of the sticks is captured and analyzed with a computer vision algorithm [17]. The system generates real-time music notation based on this analysis for instrumental musicians who can sight-read—that is, they perform from the generated sheet music as they read it, in real time. This is similar to the approach of audience response systems—that is audience voting—the outcome of which is used to shape a performer's play indirectly and dynamically [53, 57]. While musicians can safeguard the musical outcome, mediating the audience's votes toward a pre-composed sonic outcome, the musical change effected by the audience may not be readily distinguishable, preventing the audience from understanding how their participation contributes to the music they hear [32].

The idea of using mobile phones to help broader populations participate in music-making have emerged with the development of smartphones [48, 52, 51, 9]. Particularly, the ubiquity of mobile phones can resolve this problem by allowing musicians to transfer musicianship to the audience, letting them literally play musical instruments using their smartphones in their palms [35, 33, 45, 8]. The main benefit of such an approach is that audience members' engagement can directly contribute to the music and they can have a clear understanding of how individual sonic outcomes from their smartphone speakers add to the music. However, the risk of low participation can result in a complete lack of sound, which can be catastrophic for the musicians. To address this concern, musicians need to effectively sustain the audience's interest in participating. In our work, we suggest that supporting social interactivity in an audience participation system would effectively engage the audience in a sustainable fashion, while letting musicians make music with the help of the audience.

Lastly, we draw ideas from crowdsourced content creation when an artifact can be collaboratively created with crowd participants. One such example is *Scribe*, a system that coordinates a crowd, potentially in a classroom or conference, to create real-time caption of speech for the people with hearing

impairment [30, 31]. On the other hand, participants can build a physical artifact; in [28], the authors presented a networked system that can guide conference attendees to build a large-scale physical architecture with modular design. The online crowd can be incorporated in various creative practices from writing [25, 26] to animation [38]. Online crowdworkers can be recruited to support content creation in near real-time; one can hire crowd workers to design graphical user interfaces prototype [29], or to even augment the digital artifact by implementing interactive behaviors in it [37, 36]. This work addresses the similar challenges of facilitating instant and easy content creation for non-expert crowd in interactive music performance.

CROWD IN C: DESIGN CONSIDERATIONS

We realize the idea of social interactivity in a participatory musical performance in *Crowd in C*, an interactive musical system for large-scale audience involvement at a concert [33]. During a performance, each participating audience member uses a web-based, interactive musical instrument application, typically from their own smartphones. The aggregate sound outcome from all participants' smartphone speakers creates a heterophonic texture, initially centered around a C major chord. The mobile phones that the audience members use are connected through the on-stage performer's laptop via a cloud service (PubNub). More details regarding the musical motivation of the music piece can be found in [33]. The challenges of large-scale participatory music can be summarized as follows:

- (C1) Lowering the barriers in instant participation for non-experts
- (C2) Sustaining an active level of participation and facilitating the collaborative creation for participants
- (C3) Shaping audience interactions into a coherent musical piece

This paper focuses on assessing the system with respect to the second challenge (C2) in particular: that is, how the social interactivity provided in the system can effectively address this challenge. The challenge (C2) is relevant to sustaining attraction (sustainer) in creative engagement [14], expanding its context from exhibition to the live performance setting. The other two challenges and the piece's musical aesthetic in relation to contemporary music are discussed more in depth in [33, 34]. To introduce how *Crowd in C* works, design considerations relevant to each challenge are discussed below.

Democratized Music Generation

The user interface for audience members in *Crowd in C* is designed primarily for ease of use. A simple and intuitive interface is essential to motivate people to immediately participate in the piece regardless of their musical backgrounds or lack thereof, especially given that they have little time to learn or prepare. In addition, thanks to the ubiquity of smartphones and the advancement of web audio, this performance requires minimal technical configuration for the audience—participants need only launch a web browser and type in a provided URL. This serves the goal of enabling an arbitrary audience to instantly participate in the performance with minimal effort.

¹<https://www.youtube.com/watch?v=ne6tB2KiZuk>

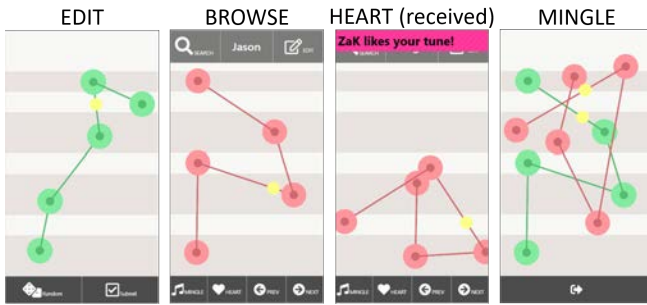


Figure 2. Screen capture of the web interface for the audience in *Crowd in C*. The system provides a simple loop-based musical instrument that anybody can play. The location of the musical notes (presented as circles) determines the pitch and timbre of the notes. Once a melody is submitted in EDIT mode, a participant can BROWSE other audience members’ melodies and interact with them by sending a HEART or playing together in MINGLE mode.

The musical instrument provided to the audience for *Crowd in C* is designed to be extremely easy to use, which is advantageous in the context of large-scale audience participation. Using the musical instrument available at a specific URL, each participant simply defines a melody on the platform (i.e., instrument) by creating a spatial pattern with five musical notes (green circles in Figure 2 - EDIT). The location of a musical note in space determines its pitch (vertical axis), timbre (horizontal axis), and the intervals between notes (the length of lines that connect notes). There is a smaller yellow circle that moves at a constant speed, triggering sound generation whenever it reaches a musical note. The composed melody is looped while a user revises the melody. Thus, a participant does not need to produce a continuous sequence of “playing gestures” (e.g., striking keys on a piano) to generate sound. More importantly, musicians do not need to worry about the overall sound being too sparse due to low audience participation. While designing a creative tool to have a low floor and high ceiling is important [43, 54], we intentionally designed the instrument to have an extremely low floor and low ceiling, as it is a musical instrument that is designed to be used only for tens of minutes, without time spent learning how to use it. This design choice contributes to the challenge of sustaining participants’ interest (C2): *if we give out a musical instrument that one can master in tens of seconds, how do we maintain users’ interest in using the instrument?* We suggest that enabling social interaction using the music instrument is a resolution to this challenge.

Social Interaction through Musical Instruments

Most importantly, the system allows audience members to interact with each other. The instrument provides a sort of ad-hoc social network where participants can interact with each other by browsing and liking (sending HEARTs to) each other’s melodies. Social interaction in *Crowd in C* uses the metaphor of online dating; the social interaction available is patterned after an online dating application (e.g., Tinder). Each interactive component is described below:

- **EDIT:** A participant can compose (or revise) a melody like a personal profile in an online dating website. Initially, participants are asked to create a screen name and compose

a melody when they launch the application. They can revise their melodies thereafter (Figure 2 - EDIT).

- **BROWSE:** Participants can browse melodies composed by others, as if they browsed people’s profiles in online dating websites. A participant can press a button labeled ‘NEXT’ to see the melody of the next participant in the queue. They can also choose to directly find a user based on their screen name using the SEARCH button (Figure 2 - BROWSE).
- **HEART:** A participant, say *Jane*, can endorse someone’s melody, say *Danae*’s, by sending a HEART. A recipient of a HEART will get a notification that reads ‘*Jane* liked your tune!’ with the sound effect of ⟨blip⟩, as shown in Figure 2 - HEART (3rd panel). If two users send each other HEARTs, both receive a notification like ‘It’s a match! *Danae* liked your tune, too’ with the sound effect of ⟨fairy wand sound⟩ on both devices.
- **MINGLE:** A participant can request a MINGLE session with another user. The user who received the request can accept or reject it. If the request is accepted, the melodies of two participants are overlaid on one screen and they can EDIT their own melodies in real time, as if chatting over a messenger or going out on a date. In this mode, one can make anticipatory revisions to match their partner’s melody either musically or visually—which brings musical intimacy, as in the example. See the upper left picture in Figure 1. The MINGLE mode terminates if either participant leaves the mode (Figure 2 - MINGLE).

Composing a profile, browsing, and endorsing online content are common interactions well-known to modern social media (e.g., Twitter, Facebook) users. We hypothesize that social interactivity will encourage participation and exchange creative ideas among the participants. We expect that the behaviors that emerge in this temporary social network help us gain insight into ways of increasing participation in large, live musical performances.

The Performer Performs with the Crowd

In *Crowd in C*, there is an on-stage musician performing the piece together with the audience members. Overall, the performer’s role is simply to perform a piece of music, explain some of the functions of the application both prior to and during the performance, and communicate with the audience. The stage performer serves as a sort of meta-performer, controlling the chord progression of the music by changing the scale to which the audience’s mobile instruments are tuned at any given moment. For example, the performer can alter the tuning of participants’ instruments from C major to E minor on the fly. The performer writes code from the stage to make these musical changes, sending JavaScript code to each connected mobile phone. The program evaluates the code it receives and makes the corresponding changes, altering the configuration (scale, key, volume, timbre) of the musical instrument. Repeating this process creates a chord progression. The performer thus does not generate any sound from the stage independently, instead controlling the high-level structure of the music by altering the configuration of the web-based musical instruments

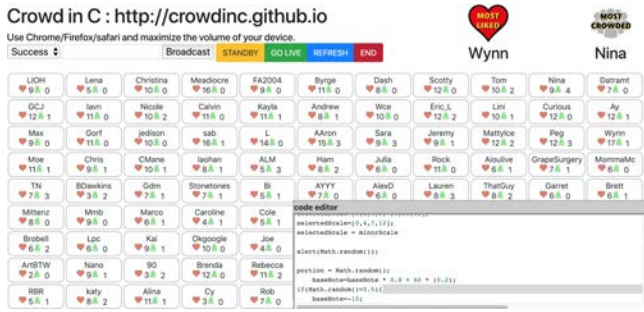


Figure 3. Screenshot of the performer’s interface from the analyzed performance. The performer live-codes in the editor (at the bottom) and sends the code text to connected smartphones in the audience to make changes in the music. In addition, the status of all connected audience members is displayed to promote the communal and competitive nature of the participatory performance.

that the audience plays on the fly. For more detail on how the performer performs and the networked structures, see [34, 10].

In addition, the on-stage performer can communicate with the audience by sending textual messages in the form of in-app notifications to the audience’s smartphones. The performer can ask the audience to make certain desirable changes to the music, or simply ask a binary question that people can answer yes or no to (e.g., “Make the melody more dense to have more notes.”, “Let’s have more high notes!”, or “Are you guys having fun?”). Pressing the ‘yes’ or ‘no’ buttons triggers a sound effect of a voice saying yes or no, making it possible to hear how others respond to the question.

The performer’s laptop is mirrored in the concert space, and shows the aggregate statistics for each individual: how many times each audience member received HEARTs and how many people are currently playing melodies, resembling a scoreboard (Figure 3). The scoreboard is designed to promote social translucence (visibility and awareness) by providing real-time information about those who are connected [15]. This information could make the audience’s own interface too complex were it presented in the web app, which participants use for a very limited amount of time. Beyond giving a holistic view of the audience, the scoreboard adds a competitive aspect to the experience—for instance, it shows the screen name of the most “liked” person—the one who received the most HEARTs—and the most “crowded” person—the one whose melody is currently being played the most at the moment.

Lastly, the performer’s computer serves as a central hub that keep tracks of and distributes all kinds of data whenever requested. The performer’s program maintains the entire queue of participants and their patterns. For example, if a user presses the NEXT button, the web app requests the screen name of the next person in the queue and their melody. The system has been used for five public performances, involving a total of over 350 participants in the wild.

METHODS

Performance

In this paper, we focus on one recent performance of *Crowd in C*, in which a musician preformed at a university computer music concert. The concert was organized by the School of

Performing Arts, and *Crowd in C* was one of many pieces performed that night. The concert was advertised to the general public through a website, social media, and a mailing list, attracting university students and local residents, including minors. We manually counted the attendees from the performance video footage and there were 87 attendees, and the data log indicates that 98 screen names were created during the performance. This discrepancy is likely explained by the presence of others (e.g., performers, staff) in the music hall and possibly people using multiple devices or browsers. This suggests that the majority, if not all, of the attendees participated in and contributed to the performance. Prior to the performance, the performer spent a few minutes explaining the concept of the piece and guiding the audience to participate in the piece using their smartphones. The URL to the application was displayed in the performance interface, which was projected on the screen so that participants could access it. Figure 3 shows screen names created by the audience on the night of the performance. The performance lasted approximately 520 seconds (or 8.7 minutes).

Data Collection and Analysis

Prior to the performance, we developed an interaction trace logger that collects complete interactive behaviors and network traffic messages. The logger maintains a separate ‘log’ channel subscribed to all the connected devices, sending a log message whenever there is a change in the program state due to user interaction or a network message. Each message contains a message type marker, the screen name of the originating user (which may be the performer), a time stamp, and any auxiliary data needed to account for the behavior. For instance, when a user sends a HEART to another user, a network message, type of which is HEART, is sent to the log channel with the auxiliary data, including the screen names of the sender and receiver. Any state transition of a user will be collected as well; if a user decides to revise their melody while browsing other users’ melodies, the state transition from BROWSE to EDIT will be logged with the screen name and time stamp. In addition, whenever a user makes changes in their melody, either in EDIT or MINGLE mode, that change is sent to the log channel.

The goal of the interaction trace logger is to enable procedural analysis of each performance of *Crowd In C*. The density of user actions can be tracked over time for the whole audience or each individual. Using the collected interaction traces, we analyzed the audience activities and examined the effects of social interaction in the system. After collecting the raw data, we aggregated it into certain relevant categories. Examples of such categories are explained below.

- **EDIT:** Changing the location of a musical note; submitting a revised melody to the centralized system—the performer’s laptop. Everyone is supposed to compose a melody after they create their screen name (initial EDIT), and they can come back to revise their melody (noninitial EDIT)
- **BROWSE:** Browsing an audience member’s melody
- **HEART:** Sending a HEART to another audience member




Figure 4. Trend Analysis of Audience Activity. The temporal analysis of each audience action is aggregated into five different types: initial EDIT, noninitial EDIT, BROWSE, HEART, and MINGLE. The result suggests that the audience has been constantly active throughout the performance.

- **MINGLE:** Requesting a MINGLE session; accepting a MINGLE request; real-time interaction in a MINGLE session—moving musical notes within a MINGLE request

In addition to the categories above, we differentiated *initial EDIT*, which was mandatory for all audience members, from *noninitial EDIT*. Noninitial EDIT activity is optional, and indicates that an audience member decided to revisit and revise their melody. The level of activity in revisiting helps us understand if audience members' intentions in interacting with the instrument are purely musical or social (e.g., a participant decided to revisit and revise a melody after seeing someone else's melody). In addition, we logged all textual messages and musical changes made by the performer with time stamps to see how the performer's intervention affects audience participation.

RESULTS

We present the results of our analysis on the data collected from the performance.

Sustaining Audience Engagement

In this section, we analyze the interaction log of the audience at the performance to assess the effects of social interaction in sustaining audience engagement.

Trend Analysis of Audience Activity

The result of trend analysis indicates that the audience was actively engaged with the participatory system over the duration of the performance. The overall activity level of the whole audience was stable throughout the performance. We present the number of activities per category in Figure 4, aggregating data points into 10-second bins. The average number of interactions per audience member was 1.43 interactions per 10 seconds. The total number of activities for the whole audience was not linearly correlated with time, which indicates that the activity level was constant during the performance ($R^2 = 0.01$).

The primary means of maintaining overall activity over time was the social elements of the system. In Figure 4, activity of types EDIT, BROWSE, HEART, and MINGLE is displayed as a stacked area chart. It is notable that initially, the entire audience started with initial EDIT activity, which was how the

system was designed for all users. Once they submitted their melodies, the audience began to focus more on the social aspects of the system—BROWSE, HEART, and MINGLE—than on modifying their own patterns—initial EDIT and noninitial EDIT, which are the two light-gray regions on the bottom in Figure 4. The trend analysis indicates that the social interactivity of the system—BROWSE, HEART, MINGLE—was a primary factor in encouraging the audience to sustain their interest in participation. While social interaction was the primary vehicle for engaging the audience with their mobile phones (BROWSE + HEART + MINGLE in Figure 4), the audience continuously revised their melodies even after their first submission (initial EDIT). The active level of non-initial EDIT showed weak linear correlation with time ($R^2 = 0.03$), which suggests that the overall level of motivation for modifying existing melodies was constant throughout the performance.

The Performer's Intervention

Further evidence that the audience was actively engaged with the performance system comes from the trace of audience responses to the performer's intervention. Using the notification system, the performer sent multiple messages to explain the system during the performance and to ask the audience to make certain changes. For example, at 210 seconds, the performer sent a message reminding the audience of the MINGLE function, through which they can interact with other audience members in real time. The graph result shows a local increase in MINGLE activity immediately after the message was sent. Additionally, the performer asked the audience at 415 seconds to send HEARTs to many patterns upon receiving a binary question at 431 seconds, generating a large number of HEARTs and matching sounds at once with a slight slowdown after the binary question at 431 seconds. Receiving HEARTs, and mutual HEARTs (match) creates short sound effects of ⟨blip⟩ and ⟨fairy magic wand sound⟩, respectively. In the recording of the actual performance, both sound effects were frequently audible after the performer's instruction. The performer's ability to make audible changes in the music by sending textual messages to the audience suggests that the audience was actively participating in the performance.

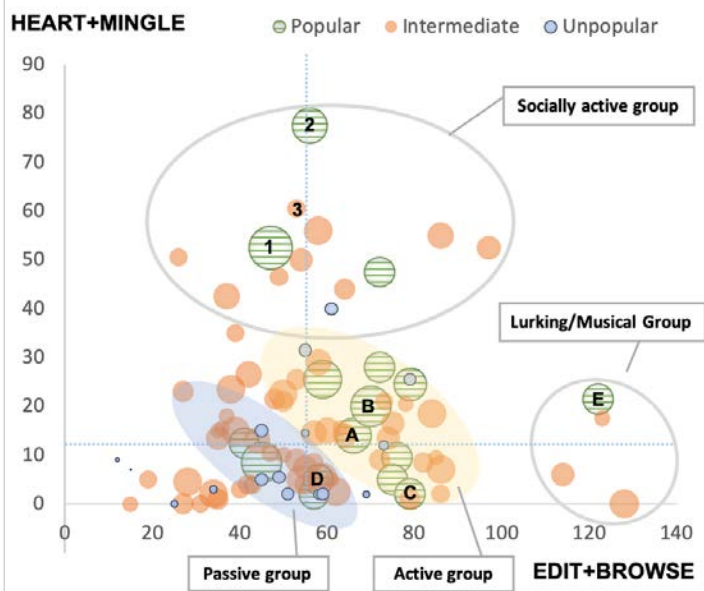


Figure 5. Emergence of different behaviors in *Crowd in C*. Each cluster of participants represents a type of behavior (socially active vs. musically active). The melodies of the users marked A-E appear in Figure 6.

Emergence of Varying Social and Musical Behaviors

There were varying behaviors from one individual to the next in social interaction. For example, the HEART functionality is a basic way to endorse someone in the system. On average, audience members sent (or received) HEARTs 8.21 times. While the standard deviation for receiving HEART is small ($\sigma = 3.48$, $\text{MAX} = 17$), sending HEARTs was driven by a small portion of people (for which $\sigma = 11.02$, $\text{MAX} = 67$)—that is, the top 20% of participants sent 62.2% of all HEARTs sent. However, the number of HEARTs sent and received showed no linear correlation ($R^2 = 0.05$).

To better understand how individuals behave differently, we plotted each audience member in Figure 5, based on how socially active they were (x -axis: HEARTs sent + MINGLE), and how musically active (or socially passive) they were (y -axis: EDIT + BROWSE). The diameter of each circle represents the number of HEARTs received, divided into three different groups — Popular (≥ 12), Intermediate ([6,11]), and Unpopular (< 6), discussed in a later section. While the majority of the audience is centered around the median (the dotted cross-hair), one group of people was more socially active than others (the cluster at the top of the graph), and another group was more focused on musical interaction with the application and passively browsing other users’ creations (the cluster near the bottom right corner of the graph). Two other clusters can be suggested near the median point: a group generally more active (*Active group*) and a group that is less active (*Passive group*). Lastly, one can see a group of people—near the bottom left corner—that were not particularly engaged with the performance system (alternatively, they may have refreshed the page and created new screen names).

The diversity of individual strategies present in the graph suggests that even though the social interactivity provided by the platform may seem overly simplistic, individuals can take var-

ious approaches to participating in the system. It is important that *Crowd in C* can accommodate broader audiences with variance in extroversion and musical expertise compared to other types of participatory models, in which audiences are asked to participate in a uniform manner. For example, in many popular music examples, audience members are often asked to clap to the beat, which can be challenging to those who have no sense of rhythm. Or, in [8], the audience was asked to tweet using a specific hashtag and some members of the audience expressed concerns that tweet messages “that make no sense outside the context of the performance” may annoy their followers, while some others did not have Twitter accounts, which limited their participation.

The Effect of Social Interactivity on Creative Engagement

In this section, we investigate the effects of social interactivity in creative engagement of the audience with the interactive music performance.

Who earns the HEARTs?

Given the intentionally limited design of the instrument, it might have been frustrating for the audience, regardless of musical background, to use the instrument in a musically meaningful way on an individual level. Therefore, people may not have been engaged with the performance in a musical sense, and instead left the sound outcome to the performer and simply focused on social interactivity. However, it seems that participants who received many HEARTs gained popularity by composing a melody in an interesting way, given the overly simplified loop-based instrument with only five notes. We musically and visually inspected the melodies of the popular group—those who are within top 20% of the audience in terms of the number of HEARTs received (received 12 or more HEARTs, $n = 16$)—and plotted as green striped circles in Figure 5. One group of participants used the instrument to create a visually meaningful pattern—readable symbols such as a letter, a shape, or the university logo (VT) created with five musical notes (Figure 6- A, B, and C, which correspond to the green circles in Figure 5-A, B, and C). One participant drew the university’s logo, which is meaningful to the local community, and thus earned 12 HEARTs (image excluded for anonymity). Another group of participants created a pattern that is musically meaningful. For example, a few participants placed five notes close to each other, or placed them in one region to play one note of a scale, so that the melody can generate more dense and unified patterns with single pitches (Figure 6-D and E, which correspond to the green circles in Figure 5-D and E). It seems that audience members clearly sent HEARTs to those who used the instrument in unexpected ways, when they could find visual or musical meaning in the limited instrument. Note that participants C and D in Figure 5 were very passive in terms of social interaction, located at the bottom of the figure. Even with the low social activity, their constant musical actions in editing their melodies elicited positive responses—(HEARTs)—from other participants.

A few people used the strategy of being highly socially active—sending many HEARTs to people. For example, the audience member who received the most HEARTs among the whole




Figure 6. Example melodies submitted by the audience. The number in the heart icon indicates the number of HEARTs that the participant received during the performance. The first five belong to the Popular group — those who received 12 or more HEARTs. Some participants decided to create visually interesting patterns (A-C), whereas others created more musically meaningful patterns (D: dense notes, E: low-pitched sound). The last two melodies were randomly selected from those who received 6 HEARTs (the third quartile).




Figure 7. The box plot of the number of EDIT triggers per group. Audience members revisited and edited their own melodies after seeing the melodies composed by members of the popular group. The average value of the Popular group was higher than that of the other two groups ($p < 0.05$).

audience—17 HEARTs—sent 34 HEARTs (Figure 5-1). However, this purely social strategy was not always effective in earning HEARTs; one participant decided to send a total of 67 HEARTs, earning only 14 HEARTs (Figure 5-2) in return. Another participant sent 55 HEARTs, gaining just 7 HEARTs themselves (Figure 5-3). We checked the melodies of these users, but observed no special patterns in the data. It seems that creativity revealed in melodic compositions had a primarily positive impact in drawing attention to other users, rather than simply being socially active.

Those Who are Popular Influence Others

Next, we investigate whether audience members musically inspired other members to edit their melodies. Towards that end, we counted how many times each melody led to a non-initial EDIT to the melody of another audience member who browsed it, which we call the number of EDIT triggers. For instance, the number of EDIT triggers for a participant we will call *John* is N if N audience members decided to revise their melodies after browsing *John's* melody. Our assumption was that if someone received many HEARTs, their melodies may be perceived as inspiring by other audience members (as shown in Figure 6 (A-E)), and a participant who saw such a melody would imme-

diately revise their own melody. We divide the whole audience into three different groups—the **Popular group**, composed of those who received 12 or more HEARTs ($n = 16$); the **Intermediate group**, composed of those who received more than 5 HEARTs but fewer than 12 HEARTs ($n = 63$); and the **Unpopular group**, composed of those who received 5 HEARTs or fewer ($n = 19$)—and compared the average number of EDIT triggers between groups. The numbers of HEARTs (12, 5) that separate the audience into three groups are determined by segmenting the group by the top 20%, middle group, and the bottom 20% in terms of the number of HEARTs received. We found that the average number of EDIT triggers for the Popular group ($\mu = 2.19, \sigma = 1.38$) was higher than those of the Intermediate group ($\mu = 1.40, \sigma = 1.20, p < 0.05$) and the Unpopular group ($\mu = 1.16, \sigma = 1.34, p < 0.05$): a two-tailed t -test (two-sample unequal variance) was statistically significant (Figure 7). The difference between the Unpopular group and the Intermediate group was not statistically significant. This strongly suggests that browsing popular melodies by others—which is relevant to the uniqueness of the pattern, as seen earlier—prompted audience members to edit their own melodies. This observation supports the idea that the social element of the system (being able to browse what others did) facilitated the audience's further musical exploration.

Musical Collaboration in MINGLE sessions

While it was used relatively less than other social elements (such as HEART and BROWSE), the MINGLE mode effectively engaged a portion of the audience in real-time creative collaboration sessions during the performance. The number of MINGLE sessions was 44, and more than half of the audience ($n = 51$) participated in one or more MINGLE sessions.

We further investigated how participants were engaged in a MINGLE session by inspecting the revision history of MINGLE sessions. Figure 8 shows selected examples that demonstrate MINGLE sessions we find interesting. For the first example (Figure 8-1), it is observed that two participants gradually made their melodies' shapes resemble each other, whereas in the second example (Figure 8-2), two participants seemed to split the canvas into two parts and spread their melodies to the sides. In the third example, two patterns started in drastically different states before taking on a harmonious appearance,

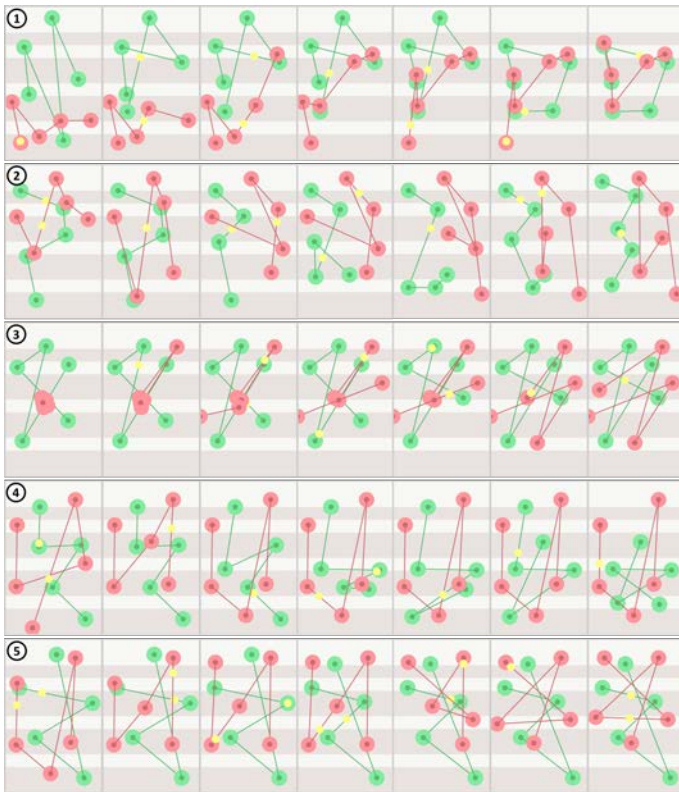


Figure 8. Curated examples of MINGLE session histories. In MINGLE sessions, participants can overlay two patterns on the screen and edit their own melodies in response to a partner’s melody and interaction.

with parallel lines and balanced placements of musical notes. The last two examples (Figure 8-4) clearly demonstrate a case of how a participant can be inspired by another audience member’s interesting approach to use the instrument. A user (Red) in example 4 saw another user making a star-shaped melody (Green). Then, the user (Red from Example 4) borrowed the same idea of making a star-shaped melody in another MINGLE session. In this case, the original user, who used Green dots in Figure 8-4, took on the role of an “influencer,” affecting their followers’ behavior. As seen above, MINGLE sessions enabled more intimate social interaction than that produced by other types of social interactions introduced in the system, positively impacting the collaborative creation process.

DISCUSSION AND LIMITATIONS

In this paper, our goal was two-fold: (RQ1) to examine how effective allowing social interaction can be in sustaining audience engagement, and (RQ2) to understand how the social interaction shapes the audience’s creative engagement with the system in a large-scale, participatory musical performance. The idea of involving social interaction is tested in *Crowd in C*, a participatory musical performance system which enables multiple layers of social interaction—performer–audience collaboration, social media-like interaction, and live collaboration. We found that the audience was generally engaged with the participatory experience: 1) the sustained social activity trend, 2) successfully coordinated musical changes with a per-

former’s intervention, and 3) the emergence of diverse patterns of participation.

A variety of social behaviors and strategies manifested themselves during the performance, similar to various types of users in today’s social media (e.g., influencers, followers, lurkers, loners). The platform’s accommodation of diverse participatory strategies—as opposed to the imposition of a monotonous strategy (e.g., voting from the beginning to the end), or an approach tuned to extroverts/musical experts (e.g. singing in public)—helps broader audiences feel engaged and connected to the participatory experience. Furthermore, this suggests that enabling social interaction in a participatory system will benefit from an audience’s understanding of existing social computing systems.

The second question we set out to answer concerns how the social interactivity shapes the audience’s creative engagement. The participatory system successfully helped audience members constantly be involved in various collaborative music making activities—composing a melody, self-expression through instruments, endorsing others’ creations, being inspired by others, and live collaborative creation.

One of the limitations of this study is that we did not match the data-driven findings with audience perception that we could learn from a follow-up study or target group interview with the participants. Other similar studies that explored participatory artworks in public settings (exhibitions, live performance) can provide methodological improvements that can be explored as a future work [14, 1, 8]. As this performance was carried out “in the wild” and was a piece of a bigger concert, we did not conduct a follow-up survey or interview with the participants for logistical reasons. Another limitation of this work is that we did not validate how the system addresses two other challenges (C1, C3), which can be improved with a similar follow-up investigation with participants.

Another aspect that we overlooked in this study is the audience’s perception of the participatory experience in relation to the performed piece of music. To shape the aggregated sonic outcome as a coherent musical piece, some intervention from the performer was inevitable, which could have distracted the audience’s social interaction with others. For example, changing the chord and the scale of the instrument (e.g., from C major to E minor) may have made the audience feel interrupted, disengaging them from their control over the instrument. Multitasking and cognitive overload in participatory systems has been studied [56], and it remains questionable how two different levels of interactions (peer-to-peer vs. large-scale participatory music) coincide and conflict with each other in the context of live content creation. Artists’ intentions can vary depending on their choice of musical aesthetic; some may be fine with any sound, as long as the audience can have an engaging musical experience; others may simply use the audience as an array of speakers [49]. Given the performance, we believe a separate study is necessary to understand the audience’s perception of the tension between audience participation as a participant and as a consumer of the artifact. We believe that a combination of qualitative research methods,

including a follow-up survey and target group interviews, will be preferable.

Future work will leverage network analysis to further explore audience engagement, and discover the effects of moderation and intervention in the social network evolution. For instance, tasks of interest include characterizing the user similarities based on their activity trends or sequences [44], and comparing the social networks that are based on different user activities (e.g., ⟨HEART⟩ and ⟨BROWSE⟩) or before/after intervention from the on-stage performer [27, 21]. Moreover, capturing user influence [47], modeling the participants' behaviors and preferences, and identifying "outlier" participants in the social network [23] (i.e., participants whose interactions are significantly different than those of the other participants) may give insights into personalized interventions from the on-stage performer to sustain audience engagement and promote collective creativity.

CONCLUSION

In this work, we sought to understand the effects of social interactivity in facilitating participatory musical performances from two specific angles: sustaining audience engagement and understanding the ways in which social interactivity influenced the audience interaction with the instrument. We tested the idea in an interactive musical performance and system, *Crowd in C*. We found that the simple, ad-hoc social network realized through the musical instruments involved the audience throughout the performance, while accommodating diverse participatory patterns and supporting various social and musical interactions among audience members. The design of including social elements in the system and its results are relevant to inform other designs that aim to facilitate participatory performances and events. We believe that our study findings will help future work better understand the design considerations of incorporating computer-mediated social interactivity for future participatory systems.

ACKNOWLEDGEMENTS

This work was generously supported by the Michigan Institute for Data Science (MIDAS) at the University of Michigan. We also thank the students in CS6724 (Spring 2019) at Virginia Tech for providing ideas on evaluation methods through an in-class discussion.

REFERENCES

- [1] Ben Bengler and Nick Bryan-Kinns. 2014. *In the Wild: Evaluating Collaborative Interactive Musical Experiences in Public Settings*. Springer International Publishing, Cham, 169–186. DOI: http://dx.doi.org/10.1007/978-3-319-04510-8_12
- [2] Tony Bergstrom, Andrew Harris, and Karrie Karahalios. 2011. Encouraging Initiative in the Classroom with Anonymous Feedback. In *Human-Computer Interaction – INTERACT 2011*, Pedro Campos, Nicholas Graham, Joaquim Jorge, Nuno Nunes, Philippe Palanque, and Marco Winckler (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 627–642.
- [3] Zafer Bilda, Ernest Edmonds, and Linda Candy. 2008. Designing for creative engagement. *Design Studies* 29, 6 (2008), 525 – 540. DOI: <http://dx.doi.org/10.1016/j.destud.2008.07.009>
- [4] Loren Carpenter and Rachel Carpenter. 1999. Audience participation. *Ars Electronica: Facing the Future* (1999), 395–396.
- [5] Teresa Cerratto-Pargman, Chiara Rossitto, and Louise Barkhuus. 2014. Understanding Audience Participation in an Interactive Theater Performance. In *Proceedings of the 8th Nordic Conference on Human-Computer Interaction: Fun, Fast, Foundational (NordiCHI '14)*. ACM, New York, NY, USA, 608–617. DOI: <http://dx.doi.org/10.1145/2639189.2641213>
- [6] Debaleena Chattopadhyay, Francesca Salvadori, Kenton O'Hara, and Sean Rintel. 2018. Beyond Presentation: Shared Slideware Control as a Resource for Collocated Collaboration. *Human-Computer Interaction* 33, 5-6 (2018), 455–498. DOI: <http://dx.doi.org/10.1080/07370024.2017.1388170>
- [7] Andrew Cross, Edward Cutrell, and William Thies. 2012. Low-cost Audience Polling Using Computer Vision. In *Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology (UIST '12)*. ACM, New York, NY, USA, 45–54. DOI: <http://dx.doi.org/10.1145/2380116.2380124>
- [8] Luke Dahl, Jorge Herrera, and Carr Wilkerson. 2011. TweetDreams : Making Music with the Audience and the World using Real-time Twitter Data. In *Proceedings of the International Conference on New Interfaces for Musical Expression*. Oslo, Norway, 272–275. http://www.nime.org/proceedings/2011/nime2011_272.pdf
- [9] Nicolas d'Alessandro, Aura Pon, Johnty Wang, David Eagle, Ehud Sharlin, and Sidney Fels. 2012. A Digital Mobile Choir: Joining Two Interfaces towards Composing and Performing Collaborative Mobile Music. In *Proceedings of the International Conference on New Interfaces for Musical Expression*. University of Michigan, Ann Arbor, Michigan. http://www.nime.org/proceedings/2012/nime2012_310.pdf
- [10] Antonio Deusany de Carvalho Junior, Sang Won Lee, and Georg Essl. 2016. Understanding Cloud Support for the Audience Participation Concert Performance of Crowd in C[loud]. In *Proceedings of the International Conference on New Interfaces for Musical Expression (2220-4806)*, Vol. 16. Queensland Conservatorium Griffith University, Brisbane, Australia, 176–181. http://www.nime.org/proceedings/2016/nime2016_paper0037.pdf
- [11] Jim DeRogatis. 2006. *Staring at Sound: The Story of the Flaming Lips*. Robson.
- [12] Honglu Du, Mary Beth Rosson, and John M. Carroll. 2012. Augmenting Classroom Participation Through Public Digital Backchannels. In *Proceedings of the 17th ACM International Conference on Supporting Group Work (GROUP '12)*. ACM, New York, NY, USA, 155–164. DOI: <http://dx.doi.org/10.1145/2389176.2389201>

- [13] Honglu Du, Mary Beth Rosson, John M. Carroll, and Craig Ganoë. 2009. I Felt Like a Contributing Member of the Class: Increasing Class Participation with Classcommons. In *Proceedings of the ACM 2009 International Conference on Supporting Group Work (GROUP '09)*. ACM, New York, NY, USA, 233–242. DOI: <http://dx.doi.org/10.1145/1531674.1531709>
- [14] Ernest Edmonds, Lizzie Muller, and Matthew Connell. 2006. On creative engagement. *Visual Communication* 5, 3 (2006), 307–322. DOI: <http://dx.doi.org/10.1177/1470357206068461>
- [15] Thomas Erickson and Wendy A. Kellogg. 2000. Social Translucence: An Approach to Designing Systems That Support Social Processes. *ACM Trans. Comput.-Hum. Interact.* 7, 1 (March 2000), 59–83. DOI: <http://dx.doi.org/10.1145/344949.345004>
- [16] C. Ailie Fraser, Tovi Grossman, and George Fitzmaurice. 2017. WeBuild: Automatically Distributing Assembly Tasks Among Collocated Workers to Improve Coordination. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, New York, NY, USA, 1817–1830. DOI: <http://dx.doi.org/10.1145/3025453.3026036>
- [17] Jason Freeman. 2008. Extreme Sight-reading, Mediated Expression, and Audience Participation: Real-time Music Notation in Live Performance. *Comput. Music J.* 32, 3 (Sept. 2008), 25–41. DOI: <http://dx.doi.org/10.1162/comj.2008.32.3.25>
- [18] Drew Harry, Eric Gordon, and Chris Schmandt. 2012. Setting the Stage for Interaction: A Tablet Application to Augment Group Discussion in a Seminar Class. In *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work (CSCW '12)*. ACM, New York, NY, USA, 1071–1080. DOI: <http://dx.doi.org/10.1145/2145204.2145364>
- [19] Drew Harry, Joshua Green, and Judith Donath. 2009. Backchan.NI: Integrating Backchannels in Physical Space. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '09)*. ACM, New York, NY, USA, 1361–1370. DOI: <http://dx.doi.org/10.1145/1518701.1518907>
- [20] Jean Hasse. 1986. Moths. *Visible Music, Euclid, OH* (1986).
- [21] Mark Heimann, Haoming Shen, Tara Safavi, and Danai Koutra. 2018. REGAL: Representation Learning-based Graph Alignment. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management (CIKM '18)*. ACM, New York, NY, USA, 117–126. DOI: <http://dx.doi.org/10.1145/3269206.3271788>
- [22] G. Jacucci, M. Wagner, I. Wagner, E. Giaccardi, M. Annunziato, N. Breyer, J. Hansen, K. Jo, S. Ossevoort, A. Perini, N. Roussel, and S. Schuricht. 2010. ParticipArt: Exploring participation in interactive art installations. In *2010 IEEE International Symposium on Mixed and Augmented Reality - Arts, Media, and Humanities*. 3–10. DOI: <http://dx.doi.org/10.1109/ISMAR-AMH.2010.5643313>
- [23] Di Jin, Aristotelis Leventidis, Haoming Shen, Ruowang Zhang, Junyue Wu, and Danai Koutra. 2017. PERSEUS-HUB: Interactive and Collective Exploration of Large-Scale Graphs. *Informatics* 4, 3 (2017). DOI: <http://dx.doi.org/10.3390/informatics4030022>
- [24] Robin H. Kay and Ann LeSage. 2009. Examining the benefits and challenges of using audience response systems: A review of the literature. *Computers & Education* 53, 3 (2009), 819 – 827. DOI: <http://dx.doi.org/10.1016/j.compedu.2009.05.001>
- [25] 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>
- [26] Aniket Kittur and Robert E. Kraut. 2008. Harnessing the Wisdom of Crowds in Wikipedia: Quality Through Coordination. In *Proceedings of the 2008 ACM Conference on Computer Supported Cooperative Work (CSCW '08)*. ACM, New York, NY, USA, 37–46. DOI: <http://dx.doi.org/10.1145/1460563.1460572>
- [27] Danai Koutra, Neil Shah, Joshua T. Vogelstein, Brian Gallagher, and Christos Faloutsos. 2016. DeltaCon: Principled Massive-Graph Similarity Function with Attribution. *ACM Trans. Knowl. Discov. Data* 10, 3, Article 28 (Feb. 2016), 43 pages. DOI: <http://dx.doi.org/10.1145/2824443>
- [28] Benjamin Lafreniere, Tovi Grossman, Fraser Anderson, Justin Matejka, Heather Kerrick, Danil Nagy, Lauren Vasey, Evan Atherton, Nicholas Beirne, Marcelo H. Coelho, Nicholas Cote, Steven Li, Andy Nogueira, Long Nguyen, Tobias Schwinn, James Stoddart, David Thomasson, Ray Wang, Thomas White, David Benjamin, Maurice Conti, Achim Menges, and George Fitzmaurice. 2016. Crowdsourced Fabrication. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology (UIST '16)*. ACM, New York, NY, USA, 15–28. DOI: <http://dx.doi.org/10.1145/2984511.2984553>
- [29] 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>
- [30] Walter S. Lasecki, Christopher Miller, Adam Sadilek, Andrew Abumoussa, Donato Borrello, Raja Kushalnagar, and Jeffrey Bigham. 2012. Real-time Captioning by Groups of Non-experts. In *Proceedings of*

the 25th Annual ACM Symposium on User Interface Software and Technology (UIST '12). ACM, New York, NY, USA, 23–34. DOI:

<http://dx.doi.org/10.1145/2380116.2380122>

- [31] Walter S. Lasecki, Christopher D. Miller, Iftexhar Naim, Raja Kushalnagar, Adam Sadilek, Daniel Gildea, and Jeffrey P. Bigham. 2017. Scribe: Deep Integration of Human and Machine Intelligence to Caption Speech in Real Time. *Commun. ACM* 60, 9 (Aug. 2017), 93–100. DOI: <http://dx.doi.org/10.1145/3068663>
- [32] Sang Won Lee. 2012. *Audience participation using mobile phones as musical instruments*. Master's thesis. Georgia Institute of Technology.
- [33] Sang Won Lee, Antonio Deusany de Carvalho Jr, and Georg Essl. 2016. Crowd in C[loud]: Audience participation music with online dating metaphor using cloud service. In *Proceedings of the Web Audio Conference*. Atlanta, Georgia.
- [34] Sang Won Lee and Gerog Essl. 2016. Live Coding the Audience Participation. In *Proceedings of International Conference on Live Coding*. Hamilton, ON, Canada.
- [35] Sang Won Lee and Jason Freeman. 2013. echobo : Audience Participation Using The Mobile Music Instrument. In *Proceedings of the International Conference on New Interfaces for Musical Expression*. Graduate School of Culture Technology, KAIST, Daejeon, Republic of Korea, 450–455. http://nime.org/proceedings/2013/nime2013_291.pdf
- [36] Sang Won Lee, Rebecca Krosnick, Sun Young Park, Brandon Keelean, Sach Vaidya, Stephanie D. O'Keefe, and Walter S. Lasecki. 2018. Exploring Real-Time Collaboration in Crowd-Powered Systems Through a UI Design Tool. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW, Article 104 (Nov. 2018), Article 104, 23 pages. DOI: <http://dx.doi.org/10.1145/3274373>
- [37] Sang Won Lee, Yujin. Zhang, Isabelle. Wong, Yang Yiwei., Stephanie. O'Keefe, and Walter S. Lasecki. 2017. SketchExpress: Remixing Animations For More Effective Crowd-Powered Prototyping Of Interactive Interfaces. In *Proceedings of the ACM Symposium on User Interface Software and Technology (UIST)*. ACM. <https://doi.org/10.1145/3126594.3126595>
- [38] Kurt Luther, Kelly Caine, Kevin Ziegler, and Amy Bruckman. 2010. Why It Works (when It Works): Success Factors in Online Creative Collaboration. In *Proceedings of the 16th ACM International Conference on Supporting Group Work (GROUP '10)*. ACM, New York, NY, USA, 1–10. DOI: <http://dx.doi.org/10.1145/1880071.1880073>
- [39] Jose Maceda. 2009. *Ugnayan*. Tzadik.
- [40] Bobby McFerrin. 2009. Bobby McFerrin Demonstrates the Power of the Pentatonic Scale. (2009). <https://www.youtube.com/watch?v=ne6tB2KiZuk>
- [41] Matti Nelimarkka, Giulio Jacucci, Antti Salovaara, Steven Dow, Kenton O'Hara, Louise Barkhuus, and Joel Fischer. 2018. Hybrid Events: Mediating Collocated Participation. In *Companion of the 2018 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW '18)*. ACM, New York, NY, USA, 455–462. DOI: <http://dx.doi.org/10.1145/3272973.3273011>
- [42] Matti Nelimarkka, Kai Kuikkaniemi, Antti Salovaara, and Giulio Jacucci. 2016. Live Participation: Augmenting Events with Audience-Performer Interaction Systems. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems (DIS '16)*. ACM, New York, NY, USA, 509–520. DOI: <http://dx.doi.org/10.1145/2901790.2901862>
- [43] Mitchel Resnick, John Maloney, Andrés Monroy-Hernández, Natalie Rusk, Evelyn Eastmond, Karen Brennan, Amon Millner, Eric Rosenbaum, Jay Silver, Brian Silverman, and Yasmin Kafai. 2009. Scratch: Programming for All. *Commun. ACM* 52, 11 (Nov. 2009), 60–67. DOI: <http://dx.doi.org/10.1145/1592761.1592779>
- [44] Tara Safavi, Chandra Sripada, and Danai Koutra. 2017. Scalable Hashing-Based Network Discovery. In *2017 IEEE International Conference on Data Mining (ICDM)*. 405–414. DOI: <http://dx.doi.org/10.1109/ICDM.2017.50>
- [45] Norbert Schnell, Sébastien Robaszekiewicz, Frederic Bevilacqua, and Diemo Schwarz. 2015. Collective Sound Checks: Exploring Intertwined Sonic and Social Affordances of Mobile Web Applications. In *Proceedings of the Ninth International Conference on Tangible, Embedded, and Embodied Interaction (TEI '15)*. ACM, New York, NY, USA, 685–690. DOI: <http://dx.doi.org/10.1145/2677199.2688808>
- [46] Jennifer G. Sheridan, Nick Bryan-Kinns, and Alice Bayliss. 2007. Encouraging Witting Participation and Performance in Digital Live Art. In *Proceedings of the 21st British HCI Group Annual Conference on People and Computers: HCI...But Not As We Know It - Volume 1 (BCS-HCI '07)*. British Computer Society, Swinton, UK, UK, 13–23. <http://dl.acm.org/citation.cfm?id=1531294.1531297>
- [47] Jimeng Sun and Jie Tang. 2011. *A Survey of Models and Algorithms for Social Influence Analysis*. Springer US, Boston, MA, 177–214. DOI: http://dx.doi.org/10.1007/978-1-4419-8462-3_7
- [48] Atau Tanaka. 2004. Mobile Music Making. In *Proceedings of the International Conference on New Interfaces for Musical Expression*. Hamamatsu, Japan, 154–156. http://www.nime.org/proceedings/2004/nime2004_154.pdf
- [49] Benjamin Taylor. 2017. A History of the Audience as a Speaker Array. In *Proceedings of the International Conference on New Interfaces for Musical Expression*. Aalborg University Copenhagen, Copenhagen, Denmark, 481–486. http://www.nime.org/proceedings/2017/nime2017_paper0091.pdf

- [50] Eric Tremblay. 2010. Educating the Mobile Generation - using personal cell phones as audience response systems in post-secondary science teaching. *Journal of Computers in Mathematics and Science Teaching* 29, 2 (April 2010), 217–227. <https://www.learntechlib.org/p/32314>
- [51] Ge Wang. 2014. Ocarina: Designing the iPhone’s Magic Flute. *Computer Music Journal* 38, 2 (2014), 8–21. DOI : http://dx.doi.org/10.1162/COMJ_a_00236
- [52] Ge Wang, Georg Essl, Deutsche Telekom, and Henri Penttinen. 2008. MoPho: Do Mobile Phones Dream of Electric Orchestras. In *Proceedings International Computer Music Conference (ICMC)*. 29.
- [53] Nathan Weitzner, Jason Freeman, Yan-Ling Chen, and Stephen Garrett. 2013. massMobile: towards a flexible framework for large-scale participatory collaborations in live performances. *Organised Sound* 18, 1 (2013), 30–42. DOI : <http://dx.doi.org/10.1017/S1355771812000222>
- [54] David Wessel and Matthew Wright. 2002. Problems and Prospects for Intimate Musical Control of Computers. *Computer Music Journal* 26, 3 (2002), 11–22. DOI : <http://dx.doi.org/10.1162/014892602320582945>
- [55] Yongmeng Wu and Nick Bryan-Kinns. 2017. Supporting Non-Musicians? Creative Engagement with Musical Interfaces. In *Proceedings of the 2017 ACM SIGCHI Conference on Creativity and Cognition (C&C '17)*. ACM, New York, NY, USA, 275–286. DOI : <http://dx.doi.org/10.1145/3059454.3059457>
- [56] Sarita Yardi. 2006. The Role of the Backchannel in Collaborative Learning Environments. In *Proceedings of the 7th International Conference on Learning Sciences (ICLS '06)*. International Society of the Learning Sciences, 852–858. <http://dl.acm.org/citation.cfm?id=1150034.1150158>
- [57] Leshao Zhang, Yongmeng Wu, and Mathieu Barthet. 2016. A Web Application for Audience Participation in Live Music Performance: The Open Symphony Use Case. In *Proceedings of the International Conference on New Interfaces for Musical Expression (2220-4806)*, Vol. 16. Queensland Conservatorium Griffith University, Brisbane, Australia, 170–175. http://www.nime.org/proceedings/2016/nime2016_paper0036.pdf