

WORLD STAGE: A CROWDSOURCING PARADIGM FOR SOCIAL / MOBILE MUSIC

Ge Wang

Jieun Oh

Spencer Salazar

Robert Hamilton

Center for Computer Research in Music and Acoustics (CCRMA)
Stanford University | Smule, Inc.

{ge, jieun5, spencer, rob}@ccrma.stanford.edu

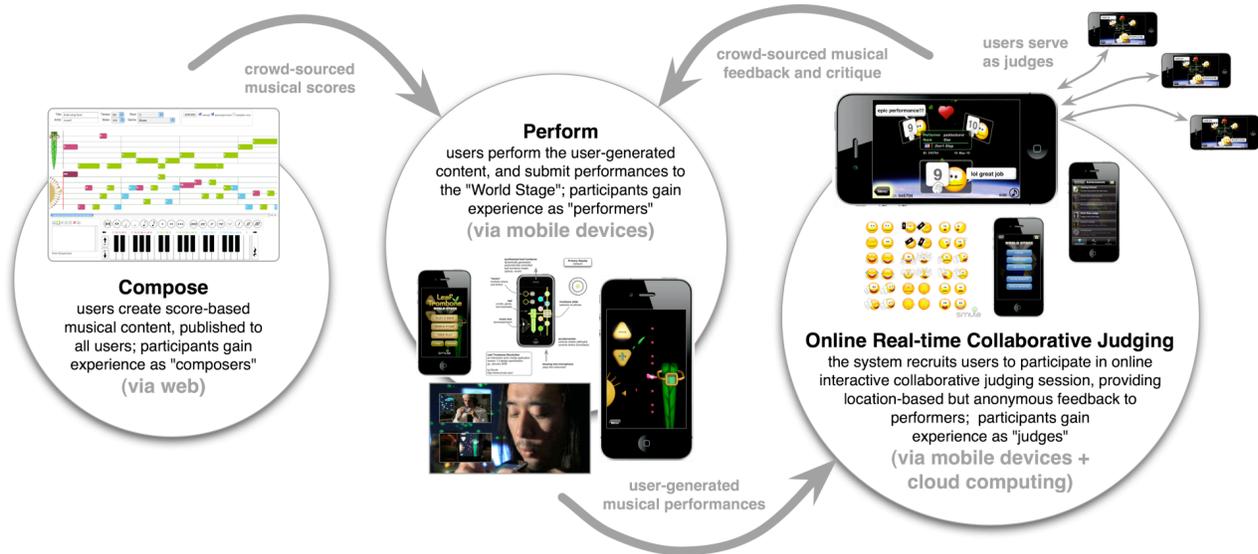


Figure 1. The mechanics of the World Stage, combining elements of crowdsourcing, cloud-computing, and human computation to create a musical ecosystem where users take on the symbiotic roles of composers, performers, and judges.

“All the world’s a stage, and all the men and women merely players...” – Shakespeare

ABSTRACT

The combination of powerful mobile devices and the connective potential of cloud-based computing are changing how, where, and when people use computers. No longer physically tethered, computers embedded in mobile phones and tablets freely roam in daily life with their human users while persistently connected to the network. At the same time, the pervasiveness of computing is enabling us to engage people to solve large-scale problems by leveraging human intelligence and the effect of large crowds. In this paper, we present a new paradigm for crowdsourced musical interactions based on mobile devices. This paradigm, which we call *World Stage*, is a pseudo-anonymous, location-aware ecosystem designed to connect hundreds of thousands of users in a social-musical game involving expressive musical performance and collaborative musical feedback. We describe the motivation and mechanics of the World Stage, and present a full implementation and case study around a commercial iPhone application: *Smule’s Leaf Trombone: World Stage*. We also present the experiential design and the technological infrastructure around the World Stage and discuss the unique social/musical possibilities it affords.

1. INTRODUCTION

The age of mobile, ubiquitous computing has begun. Powerful computers embedded in networked, mobile personal devices are transforming how, where, and when people use computing technologies - with the most profound trend perhaps being that of computers becoming increasingly pervasive and invisible (to our awareness). No longer physically, mobile phones and tablet computers are free to roam in everyday life with their human users while persistently connected to the network and aware of their location.

This combination of powerful mobile devices and the connective potential of cloud-based computing are changing the way people relate to and use computing. This paradigm shift is not only technological in nature – it is also psychological. Computing is becoming less deliberate, less static, and increasingly ad hoc, on the fly, and context-sensitive to physical interaction, location, and social relationships. As computing paradigms shift, opportunities arise for exploring new musical interactions that take advantage of the scale and pervasiveness of new technology, leading to musical experiences that are not only “mobile”, but also were perhaps never possible before.

There is another side that motivated this work. As pervasive computing is changing the way people use technology, new possibilities arise in leveraging people to solve tasks that are traditionally hard for computers,

but easy for humans. For example, the notions around Human Computation [14], Game With A Purpose [13], and cloud-based services like Amazon Mechanical Turk [1] place humans into the problem solving loop, and position technology to take advantage of human intelligence, expertise, and judgment (for example, labeling and identifying objects in photos, transcribing audio/music, researching data details). Given the undeniably human aspects of music (aesthetic, emotion, intent, expressiveness, etc.), might we not explore how technology can leverage our musical intuition and judgment for entirely new musical interactions?

1.1. What is the World Stage?

The World Stage (Figure 1) provides a massively social, game-like musical experience, leveraging the ubiquitous nature of mobile devices and ideas around human computation. It provides a platform where users can create musical content, perform using a mobile phone-based instrument, and present that performance to the greater community. As part of a social game, these users are then recruited to serve as juries in online, real-time judging sessions, giving feedback and ratings to the performers. This creates an ecosystem where a participant can take on the role of performer, judge, observer, and/or composer. The musical content, performances, and critique are all crowdsourced – that is, they originate from a large community of users. This paradigm is fully implemented and deployed in an iPhone application, Smule's *Leaf Trombone: World Stage*, which currently has a base of more than 500,000 users. We will use *Leaf Trombone: World Stage* as the primary case study to show the World Stage at work in practice, and as a way to evaluate its possibilities.

1.2. Why the World Stage?

A more important question than “what is it?” is perhaps “why?” The World Stage was conceptualized and created for one purpose: leverage the collective intelligence and judgment of the crowd to provide musical critique and feedback for user-generated performances.

World Stage was motivated by the realization that it is difficult for computers to give “deep” feedback on a musical performance, especially when it involves attributes such as expressiveness, virtuosity, musicality, and potentially unexpected qualities such as surprise and irony (e.g., it seems difficult for a computer to decide whether a performance is “funny” on some level). Yet, it is easy for a human to grasp and evaluate these features, even if the human is not a trained musician. We can recognize and appreciate virtuosity without actually being virtuosos. We are able to judge whether we like or dislike a piece of music, even if we cannot quite say why.

There is a certain universality that pervades our appreciation, understanding, and judgment of music – something that is very *human*. In a sense, the World Stage is an attempt to use technology to put our human

judgment and aesthetics into the experience, enabling new large-scale social-musical interactions.

2. RELATED WORK

2.1. Mobile Music

Mobile music-making is on the rise, accompanying the proliferation of personal mobile devices. Gaye, Holmquist, Behrendt, and Tanaka described in 2006 how mobile music can enable novel forms of musical experience by taking advantage of changes in social and geographical context [5]. Many new instruments and musical experiences have been created in recent years using mobile phones. Tanaka and Gemeinboeck pioneered locative music and media artworks [10, 11, 12]. Greg Schiemer's Pocket Gamelan [9] is one of the first examples of a mobile phone being used as a physical musical instrument. For a more complete description of related works in mobile music, see Wang, Essl, and Penttinen (2008) [17].

More recently, structured Mobile Phone ensembles have been founded across the world to explore physical interaction techniques that smartphones afford, leveraging especially their onboard sensors, network, and audio capabilities [17, 7]. In the past year, we have seen development of an API (Mobile Music Toolkit) that promotes and facilitates creating musical instruments and experiences on mobile devices [2], and a meta-environment (urMus) for designing sound and media synthesis systems on smartphones [3].

2.2. Large Group Participation

A related concept to mobile music is social connectivity and group interactions. Mobile phones can facilitate performances based on group collaboration by connecting performers in a wireless network. For instance, CaMus² (2007) by Rohs and Essl uses built-in cameras on mobile devices and Bluetooth technology to allow multiple phones to collaborate in a performance [8, 4]. Mobile devices can also enable interactions between people who may not be present in the same physical location. The iPhone's *Ocarina* [16] is perhaps the first instrument in history that allows its players to hear one another across the globe, connecting more than five million users around the world. Another example of mobile music bringing together people from geographically diverse locations is *Converge*, a composition that relies on a group of participants to capture sounds and pictures from moments in their daily lives [6]. The World Stage brings these ideas in social mobile music to a new scale and level of interactivity.

3. WORLD STAGE CONCEPT

This section chronicles the design process through which the World Stage was developed. While the final system ended up as a full-fledged realization of the World Stage, we actually did not start with the intention to build such a large-scale, crowdsourced social musical

experience. It grew out of necessity as we worked on *Leaf Trombone*.

3.1. Designing *Leaf Trombone*

Leaf Trombone was first conceived (without any notions of World Stage) in the months after the creation of *Ocarina* [16], a breath-controlled, flute-like instrument designed for the iPhone. Users blow into the microphone of the iPhone, while using the multi-touch screen to control pitch and the accelerometer to control vibrato. *Ocarina* was also designed to be a social musical experience, providing a globe visualization that allows *Ocarina* users to listen to each play around the world. A web forum was provided where users can create musical scores in the form of *Ocarina* tablature. This forum, while immensely popular, was entirely separate from the iPhone *Ocarina*.

Learning from *Ocarina*, two motivations in building a new musical instrument emerged: (1) to capture the physicality of playing the iPhone as a instrument while integrating the musical content directly into the application experience, and (2) to remain fully expressive while engaging new audiences (including novice musicians) with game-like elements. What resulted was a whimsical wind instrument we called the *Leaf Trombone*.

As a playable instrument, the *Leaf Trombone* consists of a leaf-like trombone slide accompanied by an animated Music Box (Figures 2, 3). Like *Ocarina*, *Leaf Trombone* utilizes breath control for note onset and amplitude enveloping, and, unlike *Ocarina*, is capable of smooth glissandi and a pitch range of three octaves (shiftable via two onscreen buttons).



Figure 2. *Leaf Trombone* instrument in action. Animated leaves hint where to position the green trombone slide and when to blow into the instrument, while expression is left to the user.

In game-play mode, players use the instrument to perform musical selections. In this mode, animated scrolling leaf markers prompt where to position the *Leaf Trombone* slide next and when to articulate via breath. This is similar in broad concept to games like *Guitar Hero* (2005) and *Rock Band* (2007), but with one key difference: *the sound generated in Leaf Trombone is always created by user interaction, and the user is free to express and embellish.* There is no pre-recorded

sound. In this context, it is helpful to think of *Guitar Hero* and *Rock Band* as games that *masquerade* as instruments, whereas *Leaf Trombone* is designed firstly to be an expressive musical instrument, augmented with game-like qualities.

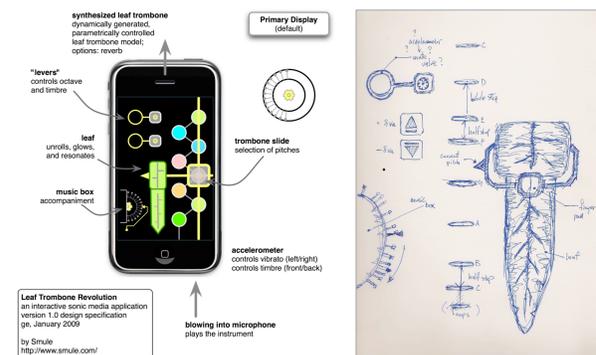


Figure 3. *Left:* an early design of *Leaf Trombone*, when it was still codenamed *Leaf Trombone Revolution*; *Right:* a later design sketch that more closely resembles the final interface.

3.2. Origin of the World Stage

At this point in the design process, a difficult problem presented itself. We wanted *Leaf Trombone* to provide some type of rating or feedback to the user, and yet we hesitated to have the software automatically judge a performance. On one hand, it seemed straightforward to program the *Leaf Trombone* to track the “accuracy” of how a piece was played (e.g., how many notes were “correctly” played). However, given the whimsical, expressive, and often goofy personality of the *Leaf Trombone* instrument, there seemed to exist more interesting criteria in judging the value of a performance – qualities like expressiveness, surprise, context, and irony (for example, using *Leaf Trombone* to play Beethoven's *Ode to Joy* would likely require a very different mindset than playing, say, Europe's *Final Countdown*). Yet, these attributes seemed fundamentally difficult for a computer to judge.

Through this reasoning process, we arrived at what felt like a logical solution. Instead of trying to use computers to solve these fundamentally hard problems, what if we leveraged the computer (and perhaps a network of computers) to connect and incorporate human intelligence and musical judgment into the system? After all, what we are after with scoring *Leaf Trombone* performances are the very human qualities surrounding a musical performance. It is easy and natural for people to grasp and evaluate these higher-level concepts (so natural it might be harder *not* to make a judgment) – we just need a way to harness this human musical intuition at a large scale. The realization of this concept soon became the *World Stage*, exploring how people inherently can make better judges in the right conditions.

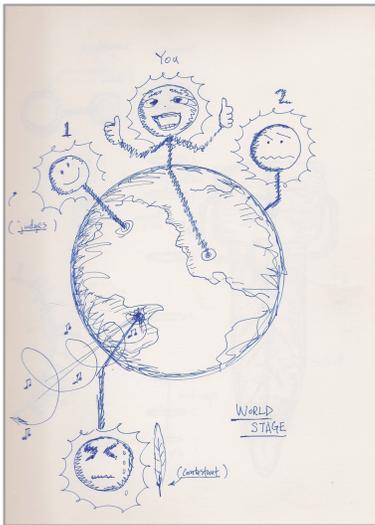


Figure 4. Initial concept drawing for World Stage. The three entities on top represent judges, and the performer is represented on the bottom. The GPS location of each user is also visualized.

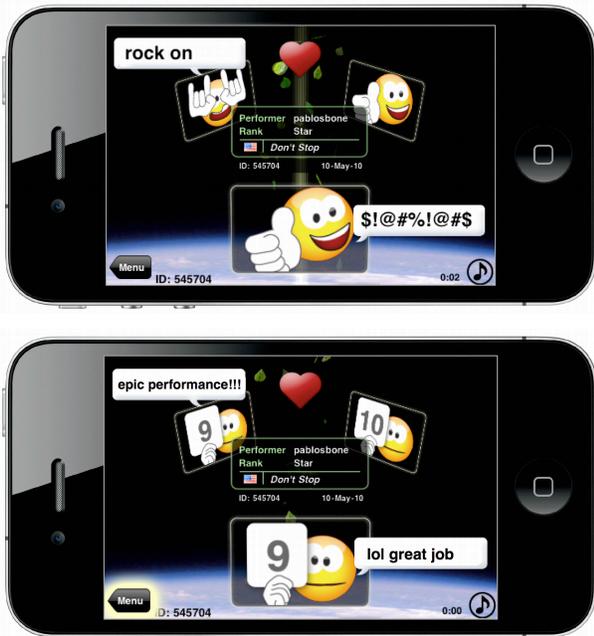


Figure 5. *Top:* a panel of three online users judges a *Leaf Trombone* performance. The judges can confer and emote in real-time with fellow judges. *Bottom:* at the end of each judging sessions, the judges each provide a score on the scale of 1 (lowest) to 10 (highest).

3.3. World Stage Mechanics

After a player performs a song, he or she can submit it to the World Stage, where it will be judged and scored by a live panel of juries consisting of other *Leaf Trombone* players from around the world (Figure 4).

In each judging session, three judges are selected from the pool of *Leaf Trombone: World Stage* users. Their task is to listen, comment, and rate a particular performance. While the performance takes place before

it is judged, the three judges participate in real-time relative to each other, and thus are able to influence one another. As the performance is played back, judges are able to express themselves using textual commentary and emoticons, ranging from enthusiastic approval to boredom and bafflement (Figures 5, 6). At the end of each performance, each judge gives a score of 1 (weak) to 10 (strong). The judging sessions are fully recorded, stored on the central *Leaf Trombone* server, and available to the performer for playback.

4. LARGE-SCALE SOCIAL MODEL

4.1. Roles in the World Stage Ecosystem

The World Stage is an ecosystem with several symbiotic roles that mutually reinforce each other. A member of this ecosystem can take on multiple roles.

- **composer** – any user can take on the role of a composer by creating and publishing musical content for others to perform. This is accomplished using a web-based interface (discussed in Section 4.5) that publishes directly to the application. A composer’s unique World Stage username is credited with the work, allowing him or her to track work’s popularity.

- **performer** – a user can play any track created by composers (in fact, the majority of users will at some point do this). A user is considered a performer after deciding to submit a performance to be judged on the World Stage. The performer is credited with that performance, and can gain feedback from judges and any observers of the performance. Over time, a performer can earn promotions in in-game rank, which is displayed in the user’s profile information as a measure of his or her overall experience.

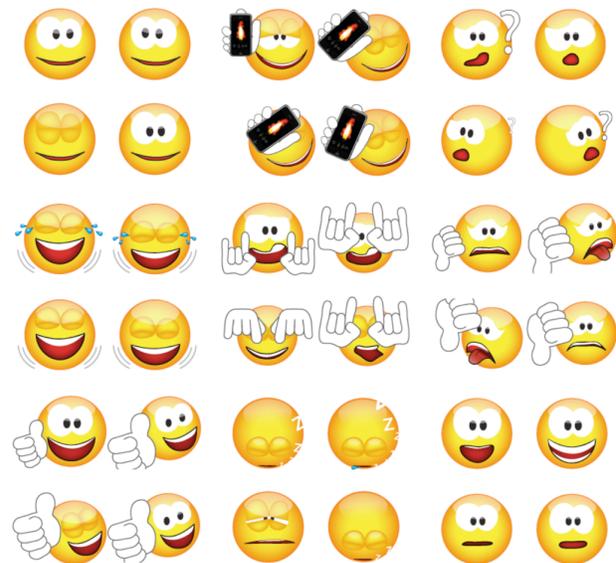


Figure 6. Judging emoticons convey sentiments about the performance.

- **judge** – judges are recruited by the World Stage to provide feedback and a final score to a performance. The World Stage strives to maintain equilibrium between the number of performers and judges. Like performers, judges can earn promotions in rank.

- **observer** – a user can observe the full recording of any existing performance that has been judged. As an observer, no direct feedback can be added, though the observer can “heart” a performance to show his or her support.

4.2. Balancing Roles

In *Leaf Trombone: World Stage*, performing for the World Stage is meant to be a recurring but special action. Users are encouraged to practice a song as much as they would like, but they have to expend a special “performance token” each time they submit a performance to the World Stage. Performance tokens can be earned over time or by serving “jury duty” (by participating as a World Stage judge). This creates a balance where users are motivated to both perform and critique. This approach works quite well – by adjusting the value of performance tokens, we can encourage judging or performances as needed, maintaining equilibrium in the system.



Figure 7. At the start of each judging sessions, the location/nationality/rank of the judges and performer are announced.

4.3. Anonymity and Location

The system regularly invites randomly chosen online users to serve as judges. Also, a user can volunteer to judge at any time. Although the system is anonymous, the following information is conveyed to fellow judges and any observers (Figure 7): the location of each judge (this is displayed in an introductory globe visualization), the nationality of each judge (visualized by a flag icon), and the rank of the judge.

4.4. Massively Multi-player Music Game

Presenting the musical experience of the World Stage as a game accomplishes several objectives. The initial perception of an experience as a “game” can lower an individual’s inhibition to start making music. Secondly, through game mechanics, we can motivate certain behaviors in an ecosystem consisting of hundreds of thousands of geographically distributed users. For

example, a token system balances the number of performances and judging sessions. Additional game elements, such as achievements and global leaderboards (Figure 8), encourage players to explore the various musical roles in the World Stage and provide a sense of accomplishment over time.



Figure 8. Left: in-game achievements for performing, judging, composing; Right: world rankings for performers.

4.5. Crowdsourcing Musical Content

Leaf Trombone: World Stage's musical content is intended to be created and renewed by its user community. Each *Leaf Trombone* song consists of a score for the *Leaf Trombone* instrument with accompaniment from the Music Box. To date, more than 5,000 user-generated scores have been published for use within the game. These scores are created using a web interface (Figure 9), which presents a notation system created specifically for the instrument. A user can compose by entering notes for the melodic trombone part and for the background accompaniment to be played on the Music Box. Scores can be previewed by listening to a MIDI version of the track, and saved to a user's online account. When it is completed to satisfaction, a user publishes the score, instantly making it available to all users of the game.

Additionally, the score is simultaneously posted to the *Leaf Trombone* online forums for discussion and viewing. Modification can be made by the original composer or other composers, and derivative works created from another user’s score are tracked as such in the central database.

5. IMPLEMENTATION

The web-based composition interface, essentially a specialized, minimal MIDI composer, is implemented using the Google Web Toolkit, a high-level API for developing rich Internet applications. When a user finishes a composition, the specified sequence of notes for the melody and accompaniment (the “score”) is serialized into a custom data format built on top of JSON. Representing scores using JSON, a widely supported, cross-platform data meta-format, allows a

potentially broad array of devices, platforms, and application developers to interact with the World Stage.

The score data and related metadata are stored in a central database, managed by MySQL. The iPhone client application dynamically and regularly updates the catalogue of songs presented to the user by querying this database over the network. When a user selects a song to play, the iPhone client retrieves the appropriate score data from the database. These and other related network calls are formalized using an HTTP middleware layer implemented in Django. Once downloaded by the client, the score data is parsed by an internal ChuckK program. The system synthesizes the Music Box accompaniment, informs the graphics subsystem what notes the user should play and when, records control data from the pitch slide and microphone “blow” input, and synthesizes the corresponding *Leaf Trombone* audio.

The recorded control input, along with the original score file, represents a complete Leaf Trombone “performance”, and if the user desires, it is uploaded to the World Stage. Due to data size concerns, we used a custom binary data format for recording performances. Input is encoded in a series of packets, each representing the blowing intensity and pitch at a specified time offset from the beginning of the performance. When offered to the World Stage, this performance data is stored in the central database and becomes available for judging.

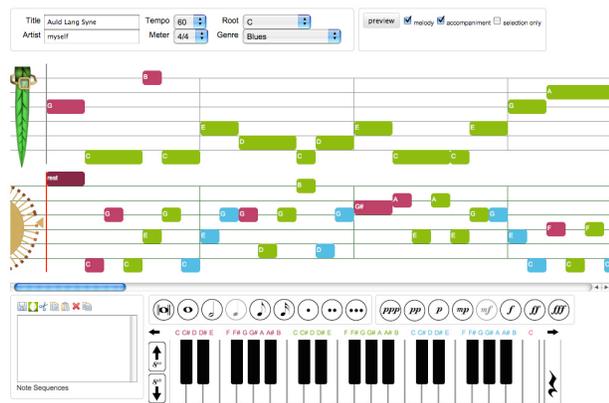


Figure 9. Composing interface for *Leaf Trombone* and the accompanying music box. This web-based application publishes directly to the application.

The World Stage judging experience, as a multi-modal real-time interaction between multiple networked human participants, requires a degree of responsiveness and statefulness difficult to attain using classical web programming environments such as PHP, Django, or Ruby on Rails. For this reason we were obliged to implement a custom stateful HTTP server to mediate between judges. When the user of the iPhone client opts to enter a judging session, the application registers its interest to the server through an initial network call. When three clients have registered in this way, the server groups them into a judging session and assigns the session a performance over which to adjudicate. If no unjudged performances are available, the server

selects from older performances that have already been judged.

Once three judges have been assembled, the session begins. The server sends each judge the performance data and score data, which is then played back through ChuckK in loose synchrony across each participant’s device. The client executes network calls to the server to register a judge’s change in emoticon or textual comment. Throughout the session, the client also polls the server for emoticon and comment updates from the other judges. The judging server timestamps and records each emoticon and comment into the database, enabling future playback of judging sessions. Upon completion of the performance, each judge offers a final numerical rating and comment, which the server records and broadcasts to each other judge.

To play back a previous judging session, the iPhone client application retrieves from the database a serialized list of the time-stamped judge updates for the session, and replays these synchronously with the recorded performance and Music Box accompaniment.

6. AFTERMATH AND MORE EXPERIMENTS

6.1. Community

Currently *Leaf Trombone: World Stage* has an estimated base of more than 500,000 users on the iPhone and iPod Touch. Users have participated in nearly 1 million judging sessions. Users have generated more than 5000 individual songs, all of which are available to other users. The most played and judged selections include renditions of *Amazing Grace*, The Beatles’ *Yesterday*, *Hallelujah*, Journey’s *Don’t Stop Believing*, *Ode to Joy*, *Pachelbel’s Canon in D*, and *Phantom of the Opera*. The most prolific composer in the World Stage single-handedly contributed 177 different songs to the community.

6.2. Visualization

To better understand and evaluate the user experience in the World Stage, we created a visualization that dynamically fetches recent judging sessions from the server and renders key features in them. The visualization was implemented as a desktop application using OpenGL and the C++ programming language.

The visualization elucidates the judging context by displaying the timestamp and latitude-longitude information of the performer and judges. The timestamp is shown relative to the present moment, such that viewers see how many minutes and seconds ago the judging session occurred (according to the records in the database tables). For displaying the location, we use reverse geo-coding to determine the city and country names corresponding to the latitude and longitude information, making the information more meaningful. Such information helps us better understand the user base, and to ask questions about emerging patterns relating the time of day, user locations (countries), and types of songs performed.

The actual performance that is being judged in a session can be triggered from within the visualizer, and re-synthesized and played back using ChucK [15]. The synchronized audio and visuals give us a clearer view of judges' behaviors. Other features of the judging session, including the average score, emoticons used, and comments given out by the three judges, are also rendered using colorful textures in such a way that the visualization serves as an artistic and informative summary of what was actually experienced by users moments ago.

6.3. Dichotomous Harmonies, a Meta-Composition

Performance data generated by World Stage users forms the basis of the musical meta-composition, *Dichotomous Harmonies*. Premiered in Milano, Italy, at the 2010 Torino-Milano Festival in a special concert event entitled *Play Your Phone!*, *Dichotomous Harmonies* combines a pre-rendered soundscape with a live improvisatory duet between a remote Trombonist (located in Bournemouth, England) and an iPad performer. At the same time, a video artist controls the World Stage visualization engine with a customized set of performance excerpts, providing a visual and conceptual backdrop to the music.

The central unifying element of *Dichotomous Harmonies* is a melodic excerpt from composer Leonard Cohen's *Hallelujah*, combining a short verse section with a single chorus. Leaf Trombone performance data for one-thousand user renditions of this excerpt were extracted from Smule's database and mixed into an 8-channel soundscape. This rendering is played back while a live performer manipulates excerpts from the same set of performances via an iPad control interface to ChucK processes running on a laptop. The improvisatory element of the iPad performance is mirrored by a networked live trombone improvisation, lending clarity to the melodic material itself while reinforcing the disembodied nature of Cohen's haunting melodic line. Combining the live trombone with the massive number of World Stage performances creates a striking dichotomy between the tangible and intangible nature of melody, performance, and presence.

7. EVALUATION AND DISCUSSION

The World Stage is a social and musical experiment at an intersection of mobile music, cloud computing and human computation. While fully implemented in *Leaf Trombone*, we believe the concept and architecture of the World Stage is applicable to other instruments and social-musical experiences.

Incorporating human intelligence into a social-musical system can be powerful, as demonstrated by abandoning computer-based judging in favor of more crowdsourced, human participation. Additionally, by adding game elements to expressive musical experiences, we can motivate more musical participation and engagement.

Another idea not to be overlooked is that anonymous interactions can be more powerful than ones where identities are known. For example, performers may be *more* likely submit performances to be judged by other users if they know the feedback is anonymous. While the validity of this point deserves more investigation, it seems clear the World Stage operates as an ecosystem without the need for traditional identities.

Finally, this experiment views crowdsourcing as marriage between technology and masses of people. We believe this is only a beginning. By using the technology to set the right conditions, we can take deeper advantage of human intelligence, judgment, and intuition in creating new musical experiences.

8. ACKNOWLEDGEMENTS

This work has been the results of many people working together, and is indebted to Jennifer Wu, David Zhu, Mattias Ljungstrom, Tricia Heath, Arnaud Berry, Jeffery Smith, Mike Earhart, Tony Ventrice, Ryan Pfenninger, Turner Kirk, Jonathan Berger, Elon Berger, and the entire Smule team. Special thanks to Clara Valenstein for analytics and analysis and to David Kerr for support and suggestions. Additional support for this research was made possible by a National Science Foundation Creative IT Grant, No. IIS-0855758.

9. REFERENCES

- [1] Amazon Mechanical Turk. <http://aws.amazon.com/mturk/> extracted January 2011.
- [2] Bryan, N. Herrera, J., Oh, J., and Wang, G. 2010. "MoMu: A Mobile Music Toolkit." In *Proceedings of the International Conference on New Interfaces for Musical Expression*. Sydney Australia.
- [3] Essl, G. and Müller, A. 2010. "Designing Mobile Musical Instruments and Environments with urMus." In *Proceedings of the International Conference on New Interfaces for Musical Expression*. Sydney, Australia.
- [4] Essl, G. and Rohs, M. 2009. "Interactivity for Mobile Music Making", *Organised Sound*, 14(2): 197-207.
- [5] Gaye, L., Holmquist, E., Behrendt, F., and Tanaka, A. 2006. "Mobile Music Technology: Report on an Emerging Community", In *Proceedings of the International Conference on New Instruments for Musical Expression*. Paris, France.
- [6] Oh, J. and Wang, G. 2011. "Audience Participation." In *Proceedings of the International Computer Music Conference*. Huddersfield, UK.

- [7] Oh, J., Herrera, H., Bryan, N., Dahl, L., and Wang, G. 2010. "Evolving the Mobile Phone Orchestra." *In Proceedings of the International Conference on New Interfaces for Musical Expression*. Sydney, Australia.
- [8] Rohs, M. and Essl, G. 2007. "Camus2 – Collaborative Music Performance with Mobile Camera Phones." *In Proceedings of the International Conference on Advances in Computer Entertainment Technology*. Salzburg, Austria.
- [9] Schiemer, G. and Havryliv, M. 2006. "Pocket Gamelan: Tuneable trajectories for flying sources in Mandala 3 and Mandala 4." *In Proceedings of the International Conference on New Interfaces for Musical Expression*. Paris, France.
- [10] Tanaka, A. 2004. "Mobile Music Making." *In Proceedings of the International Conference on New Interfaces for Musical Expression*, 154–156. Hamamatsu, Japan.
- [11] Tanaka, A. and Gemeinboeck, P. 2006b. "A framework for spatial interaction in locative media." *In Proceedings of the International Conference on New Interfaces for Musical Expression*. Paris, France.
- [12] Tanaka, A. and Gemeinboeck, P. *net d'arrive*. Project web page, 2006. <http://www.ataut.net/site/Net-Derive> Accessed January 20, 2011.
- [13] Von Ahn, L. and Dabbish, L. 2004. "Labeling Images with a Computer Game." *In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery. pp. 319–326.
- [14] Von Ahn, L. 2005. *Human Computation*. PhD Thesis. Carnegie Mellon University.
- [15] Wang, G. 2008. *The ChucK Audio Programming Language: A Strongly-timed, On-the-fly Environ/mentality*. PhD Thesis. Princeton University.
- [16] Wang, G. 2009. "Designing Smule's iPhone Ocarina." *In Proceedings of the International Conference on New Interfaces for Musical Expression*, Pittsburgh, USA.
- [17] Wang, G., Essl, G., and Penttinen, H. 2008. "Do Mobile Phones Dream of Electric Orchestras?" *In Proceedings of the International Computer Music Conference*. Belfast, Ireland.