Lagu Hantu: Creating computationally-mediated ensembles across time

Anastasha Rachel Gunawan*
gunawan@stanford.edu
Stanford University
Stanford, USA

Matthew Caren*
mcaren@stanford.edu
Stanford University
Stanford, USA

Figure 1: The *Lagu Hantu* installation.

Abstract

Lagu Hantu is a large-scale interactive instrument inspired by the Indonesian gamelan tradition, which invites people to create sound in collaboration with both each other and past participants by using distributed tactile interfaces. When the instrument is played, a program summons fragments of past participants' performances by recalling how people accompanied similar gestures in the past. We describe the design and technical implementation of the instrument, the underlying system that allows participants to collaborate across time, and the instrument's use in a public installation.

Keywords

cross-temporal collaboration, Indonesian gamelan, sound installation

1 Introduction

Music in Indonesia has long served as a bridge across temporal boundaries, enabling communities to maintain ancestral relationships, connections with past performances, and collective memory [14]. Rich practices of sound-making and craft underpin the Indonesian gamelan tradition, where music-making is inseparable from broader spiritual and social structures. Fundamental to these traditions are a cyclical conception of time and persisting relationships between the living and spirits of the past, both of which manifest throughout daily life [21], ritual and language [3], and musical practice [20].

This temporal dimension of music-making becomes particularly salient in the gamelan practice, where musical meaning emerges from accumulated gestures, inherited patterns, and the cyclical interplay of past and present performers [5, 14]. Technology offers an opportunity to both embody and extend these traditions, building on the material forms and interactive practices of gamelan as well as culturally-imbued temporal and relational principles. This perspective offers a fruitful lens for the development of NIMEs—not only for its cultural and sociological depth, but also because it inspires us to explore how musical instruments might facilitate communities across time, allowing participants who never intentionally coincide to have the experience of performing together. Digital technology offers new possibilities for embodying these cross-temporal relationships in musical interfaces and experiences.

Here, we present *Lagu Hantu* (Bahasa Indonesia: "ghost song") (Figure 1), a large-scale interactive musical instrument installation that enables participants to form asynchronous ensembles across time. As the instrument is engaged with, it accumulates a growing archive of past interactions. When new interactions occur, the system retrieves past ensemble performances based on gesture similarity, inviting participants into collective musical moments that unfolded long before their arrival.

*Both authors contributed equally to this research.

NIME '26, June 23–26, 2026, London, UK

2 Background

2.1 Related Systems

Interactive sound installations have long offered a fertile context for the development of new musical interfaces, including many examples from within the NIME community [12]. In this spirit, the *Lagu Hantu* installation inspired the development of a system that explores the interplay between three themes: musical collaboration across time, excitation of physical objects as an intuitive interface for digital instruments, and the Indonesian gamelan tradition. Although the specific intersection of these three ideas has not been well-explored, significant bodies of work have focused on them independently. We briefly review past work surrounding each.

Digital technology has become a widely-used tool in instruments and installations to connect people across space and time. Systems built to connect people across distances use computer networks [9, 17], virtual reality [29], and predictive modeling [15, 24] to unlock the experience of live co-located collaboration even when other participants are not physically present. Although less-explored in NIMEs, there are also many systems that use technology to invite people to connect across time, such as installations that embody current participants beside past participants (e.g. [19]). The MirrorFugue project explored one realization of this concept in a musical instrument, combining a player piano with projected video to conjure the image and sound of a particular past performance [33, 34]. On a surface level, many recording and sample-based systems for music also enable asynchronous musical collaboration, from multitrack tape to live looping and sample remixing systems [2, 11]. However, most musical systems that facilitate more interactive and immediate collaboration across time often do so incidentally rather than explicitly targeting the sensation of engaging with other participants whose interactions exist in the past. Many of these systems respond to events in the present by invoking and recombining prior performances. Systems for machine improvisation like OMax Brothers [1] and the Continuator [26] construct models of previous performances and recombine them to accompany live performers, effectively allowing past performances to participate in current musical interactions. In event-triggering systems like Antescofo, stored performances or compositional structures condition the timing and behavior of live musical processes [8]. More recently, AI systems trained on datasets of past performances extend these ideas by learning representations of corpora to generate new material in real-time [6, 7]. Although these systems are rarely framed as time-spanning collaborations, they nonetheless mediate interactions between present and past performers, positioning historical material as an active participant in live performance.

Physical objects augmented with acoustic sensing have proven to be a powerful interaction modality for new musical instruments, as they are able to intimately capture nuances of a performer's touch, strike, and movement [25]. Instruments like the BladeAxe [22], Mogees [23], and hybrid acoustic-digital instruments [10, 28] combine digital signal processing with material affordances to create responsive, intuitive interfaces.

Little engineering attention has been dedicated to NIMEs based on gamelan, so we review a handful of documented examples. The Gamelan Elektrika is a full set of seven MIDI instruments modeled after a Balinese Gong Kebyar that use a combination of force-sensitive resistors, piezos, and capacitive sensing to capture instrument interactions [27]. The Gamelan Elektrika, as well as the portable Gamelan Sampul interface [32], aim to both extend the sonic capabilities of gamelan-style interfaces and make them more accessible. The Gamelatron [13] is a robotic gamelan ensemble that uses solenoid actuators to excite gongs, and the Robotic Reyong [16] uses a similar approach to actuate a *reyong* instrument. Much of this prior work is built around the gamelan instruments themselves, with less attention given

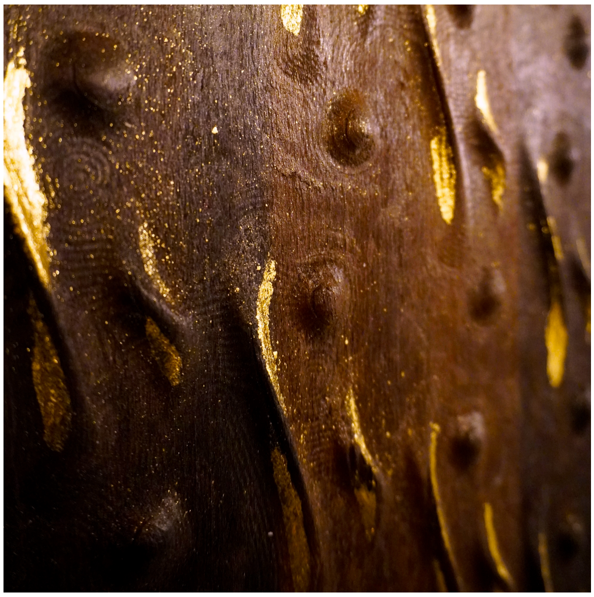

Figure 2: Detail of one panel from the Lagu Hantu interface, featuring a relief carving of a *parang batik* pattern.

to extending the ritual and temporal contexts that shape how the music is situated socially and culturally in Indonesia. *Lagu Hantu* explores a new instrument created to surface relationships with non-linear, cyclical notions of time and mediate connections between living participants and past presences.

2.2 Gamelan Traditions and Temporal Frameworks

In many Indonesian regions, life is understood as a cyclical unfolding—a continuous, circular stage where individuals grow into predestined roles [21]. Ritual practice reinforces this temporal orientation—ceremonies like *Langen Bronto* are performed through cyclic repetition, often without clear demarcations between beginning and ending, embodying an experience of time that loops upon itself [3]. This worldview appears even in linguistic structure, as Balinese, Javanese, and Indonesian are tenseless languages with no inherently implied indications of time [3]. Musical structures in gamelan similarly manifest this through their colotomic forms that return to a central gong cycle, creating patterns of recurrence [5, 20].

Underlying this cyclical worldview is the thin veil between the living and ancestral spirits. In Javanese culture, the boundary between past and present remains permeable and maintained through ritual and artistic practice [31]. Gamelan performance itself, often accompanying a *wayang* (shadow puppet) performance, functions as a medium for connecting with ancestral spirits [31], with instruments often regarded as sacred objects housing spiritual power [4]. Music meaning-making in Javanese culture, then, extends beyond individual expression and into an act of temporal communion, where present performers join in dialogue with those who came before.

Lagu Hantu is grounded in these cultural principles by exploring new ways to protect and celebrate these permeable boundaries. We honor the Indonesian traditions that inspire this work while exploring how new interfaces might materialize relationships with the past beyond static history and into persisting communion.

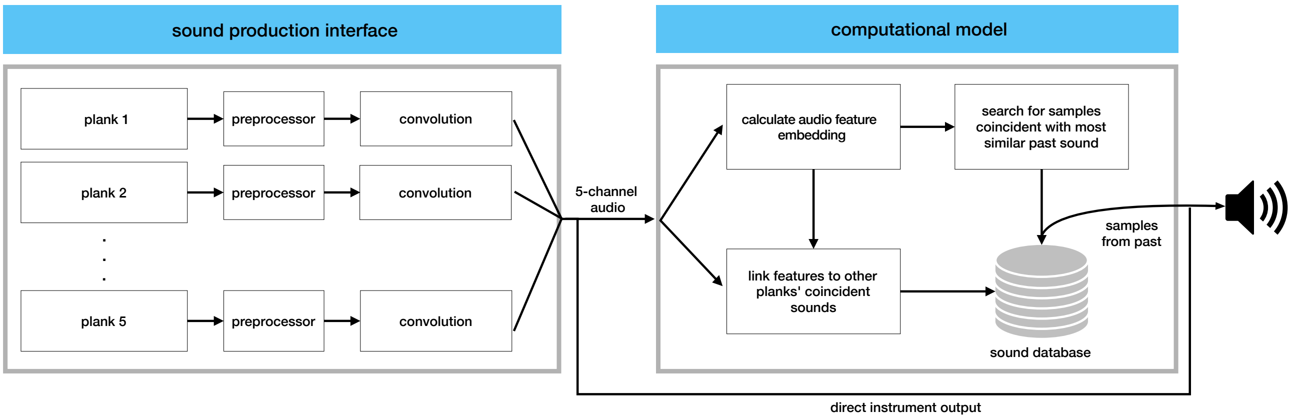

Figure 3: The *Lagu Hantu* system architecture, consisting of the sound production interface (left) and the computational model (right).

3 *Lagu Hantu*

The *Lagu Hantu* system (Figure 3) consists of two high-level components: the sound production interface and the computational model. The sound production interface consists of a set of wooden panels whose vibrations are sensed, preprocessed, and convolved with impulse responses to render sound in real-time. The computational model consists of a system that analyzes interactions with the panels and records both the audio samples and their features to a database, along with a program that uses the extracted features of current interactions to retrieve and play back accordant accompanying performances from the past.

3.1 Sound Production Interface

The physical interface consists of five hardwood panels arranged in a wide semicircle, such that participants walk from one panel to another to interact with the instrument. The installation is never silent: a continuous background soundbed plays throughout the space through a distributed array of loudspeakers.

Piezo sensors are built into the reverse side of each panel. When a participant interacts with a panel, the piezo records the resulting vibration and wires the signal through an audio interface to a central computer in a neighboring room. The recorded signals are preprocessed using a bank of bandpass filters and an envelope-controlled transient shaper to isolate participants' interactions from incidental noise. The resulting signals are then convolved with various impulse responses of gamelan gongs, with each panel being mapped to a different pitch.

The instrument is designed such that no participant ever plays unaccompanied. When multiple participants interact with different panels simultaneously, the system both records their concurrent sounds as a linked ensemble snapshot and sonifies their interactions in real time. When a single participant interacts with a panel, the system retrieves a past ensemble that included a prior participant who made a similar gesture on that panel, and plays those past sounds back through the speaker system alongside the live interaction. In this way, a solitary participant is always accompanied by an ensemble from the past. Crucially, past ensemble records are only played back when a single participant is active; when multiple participants are present and playing together, the live ensemble takes precedence and no archival playback is triggered.

The impulse responses are tuned according to the *pelog* scale, one of two principal tuning systems used in Indonesian gamelan music. The complete *pelog* scale consists of seven tones with unequal intervals. In practice, many compositions use *pelog selisir*, a five-tone subset from the available seven [30]. *Lagu Hantu* uses one such five-tone subset, tuned to approximately F#, G, A, C#, and D (i.e. roughly corresponding to scale degrees 1–♭2–♭3–5–♭6 in Western notation, though these representations do not authentically capture the non-tempered intervals of *pelog* tuning [14, 30]).

Note that unlike in similar digital percussion instruments that synthesize sound by convolving contact microphone outputs with impulse responses (e.g. [28]), we do not process the input signal to remove resonances of the original physical object, since the panels themselves were intentionally formed to have specific timbral qualities.

The panels feature different combinations of carvings and illustrations to invite a diversity of physical gestures—participants are encouraged to explore a variety of interactions, including tapping, rubbing, striking, and tracing the wooden surfaces. They are carved and gold-leafed in *parang* (Figure 2), a traditional batik pattern from ancient wax-resist textile dyeing practices. In Indonesian craft traditions, batik is not only decorative; its motifs carry symbolic meaning. The ancient *parang* form, one of the oldest batik motifs in the region, signifies continuity and unbroken flow, resonating with the cyclical temporalities that shape gamelan music and broader communal life [18]. To create the reliefs, the *parang* pattern was hand-traced and transformed into a 3D model, which was used to machine the custom wood panels using a CNC router.

3.2 Computational Model

The live signals recorded from the piezos built into each panel are continuously recorded and each segmented into three-second samples. A Python program extracts audio features from each sample in real-time to produce a corresponding vector embedding representation. This embedding consists of four scalar features for each interaction:

- **Loudness**: root mean square (RMS) energy of the sample.

- **Loudness flatness**: ratio of the geometric mean to the arithmetic mean of RMS loudness of windowed frames of the sample.
- **Spectral centroid**: intensity-weighted average frequency of the sample's frequency-domain spectrum.
- **Spectral flatness**: ratio of the geometric mean to the arithmetic mean of the sample's frequency-domain spectrum.

Each sample's feature vector is written to a database along with pointers to the coincident audio samples that were recorded from the other panels during the same three-second time window. This links each feature vector associated with an interaction at a moment in time to a multichannel sound snapshot of an ensemble moment.

When someone interacts with a panel, the system calculates audio features from the interaction in real-time, which are then compared to all existing sample embeddings from that same panel to find the nearest neighbor in feature space. For each past interaction, we compute its similarity to the current interaction in the audio feature embedding space using weighted Euclidean distance with hand-tuned weights for each feature dimension. We identify the entry with the smallest distance to the current interaction, which is interpreted as the most similar past interaction with that panel.

Once the most similar past interaction is identified, the system retrieves the sounds that were being played concurrently on the other panels at the time of that past interaction. These retrieved samples are routed through the same convolution process as the live signal to render it into tonal sound and played back alongside the participant's live interaction, allowing the participant to be accompanied by an accordant ensemble from the past.

We implemented the nearest-neighbor sample retrieval from the database using simple linear search, which still offers sub-millisecond latency per query for a database of several thousand samples.

To avoid conflicts between real-time collaborative performance and playback of past performances, samples of past performances are only audible when just a single panel is being interacted with. Because the system plays back sounds that accompany a similar past interaction—rather than the interaction's own sound—recording and feature extraction occur only when multiple panels are active simultaneously as these are the moments that would yield audible accompanying sounds. Interaction activity is detected using loudness thresholds on the input from each panel.

Note that because we are explicitly targeting the experience of playing in an ensemble alongside past performers, we intentionally avoid extensive recombination of prior performances or the creation of a generative model to accompany participants' interactions. Our goal is not to produce new material, but to faithfully echo past interactions to encourage co-performance across time. We do not attempt to evaluate or discriminate the musicality or intentionality of different participants' interactions; all past performances (any detected interactions with the instrument) are treated as being equally worthy of future playback.

4 Discussion

4.1 Public Installation

Lagu Hantu was installed in a public gallery space, where visitors were invited to freely listen to, approach, and interact with the instrument (Figure 4). The installation was presented alongside brief suggestions of various interactions they might use to explore the interface, and audio was spatialized using a multichannel speaker system. In addition to the sounds produced by the instrument itself, a background soundbed was played throughout the evening. We cataloged installation participants' behaviors through naturalistic observation as an informal but valuable exploration of the interactions and experiences that emerged *in situ*.

Figure 4: People interacting with the *Lagu Hantu* installation interface.

Interaction behaviors varied considerably across participants, and people made a wide variety of sounds (Figure 4). Some approached the panels cautiously, lightly tapping or brushing the surfaces as they tested their responsiveness. Others quickly adopted more percussive or gestural interactions, exploring dynamic variation in rhythm and texture. Interaction times also spanned a wide spectrum: while some participants explored brief touches lasting under a minute, others sustained deep engagement over extended minutes-long periods.

We especially observed that many participants, upon becoming familiar with how the instrument synthesized and retrieved sounds in real-time, oscillated between periods of sustained interaction and active listening to the instrument's output.

Although there was no prescribed symbolism or ceremony built into the installation, several participants noted aspects of ritual, describing their experiences as feeling "religious," "divine," "giving peace," and "ancient."

4.2 Future Work

The *Lagu Hantu* system remains in active development. In future iterations, we plan to expand both the physical interface and the computational model that mediates sonic collaboration across time. High-quality audiovisual documentation of the installation is currently being compiled, which we hope to release publicly in the near future.

The sound production interface could be extended to enable control over additional sonic axes like pitch, allowing for deeper—though also more demanding—participant interactions with the instrument. Expanding the scale and number of interactive panels might also open up opportunities to capture richer inter-participant dynamics.

The computational model could also be extended to encompass other notions of similarity beyond timbre like rhythm, or even

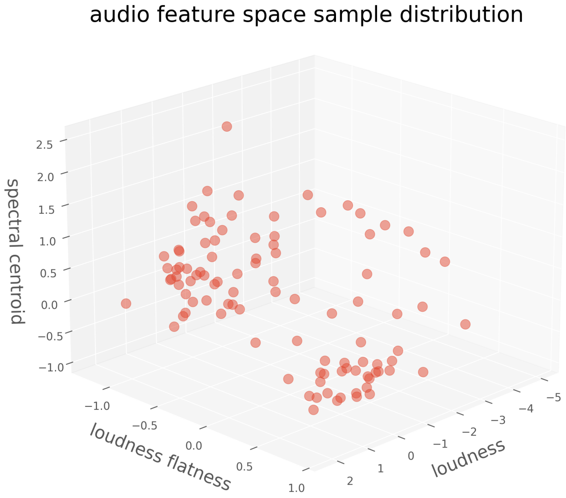

Figure 5: Distribution of interface interaction samples from one installation session, visualized in audio feature space (three randomly-selected audio feature dimensions plotted). Each point represents one 3-second sample.

inferred higher-level features like energy or narrative. Beyond the current algorithm that searches for past accompaniments of similar gestures, the system could intentionally favor contrasting gestures—for example, pairing high pitches with low pitches, or dense textures with sparse ones. Future instantiations could also explore ways of connecting participants to not only past performances, but also directly to past participants through contextual metadata or participant-generated annotations. As the installation grows to contain interconnected performances and collaborations across times and geographic regions, it will become an incrementally deeper repository of collective sonic memory.

This work situates itself within a broader line of inquiry into how musical interfaces can connect people across time. Beyond current focuses in the literature on using complex generative models trained on large datasets to accompany live performances, we hope to suggest an alternative perspective for live computationally-generated accompaniment where the source material is explicit and foregrounded.

5 Conclusion

Lagu Hantu is a system that uses a novel musical interface to not only invite intuitive sound-making interactions, but also to facilitate tangible connections across time. By grounding the interface in rich ancient tradition, the work celebrates technology as a medium for new ritual, connection, and meaning-making.

Acknowledgments

AG is supported by the Stanford Hasso Plattner Institute of Design. MC is supported by the Hertz Foundation and the Steve Jobs Archive.

6 Ethical Standards

No conflicts of interest were identified.

References

[1] Gérard Assayag, Georges Bloch, Marc Chemillier, Arshia Cont, and Shlomo Dubnov. 2006. OMax brothers: a dynamic topology of agents for improvization learning. In *Proceedings of the 1st ACM workshop on Audio and Music Computing Multimedia*. 125–132.

[2] Jeronimo Barbosa, Marcelo Mortensen Wanderley, and Stéphane Huot. 2017. Exploring playfulness in NIME design: The case of live looping tools. In *Proceedings of the 2017 New Interfaces for Musical Expression Conference (NIME '17)*. 87–92.

[3] Judith Becker. 1981. Hindu-Buddhist time in Javanese gamelan music. In *The Study of Time IV: Papers from the Fourth Conference of the International Society for the Study of Time*. Springer, 161–172.

[4] Judith Becker. 1988. Earth, fire, śakti, and the Javanese gamelan. *Ethnomusicology* (1988), 385–391.

[5] Judith Becker and Alton Becker. 1981. A musical icon: Power and meaning in Javanese gamelan music. In *The Sign in Music and Literature*. University of Texas Press, 203–215.

[6] Christodoulos Benetatos, Joseph VanderStel, and Zhiyao Duan. 2020. BachDuet: A deep learning system for human-machine counterpoint improvisation. In *Proceedings of the 2020 New Interfaces for Musical Expression Conference (NIME '20)*.

[7] Lancelot Blanchard, Perry Naseck, Stephen Brade, Kimaya Lecamwasam, Jordan Rudess, Cheng-Zhi Anna Huang, and Joseph Paradiso. 2025. The jam_bot, a real-time system for collaborative free improvisation with music language models. In *Proceedings of the 2025 International Society for Music Information Retrieval Conference (ISMIR)*.

[8] Arshia Cont. 2008. ANTESCOFO: Anticipatory Synchronization and Control of Interactive Parameters in Computer Music. In *Proceedings of the 2008 International Computer Music Conference (ICMC)*. 33–40.

[9] Juan-Pablo Cáceres and Chris Chafe. 2010. JackTrip: Under the Hood of an Engine for Network Audio. *Journal of New Music Research* 39, 3 (2010), 183–187.

[10] Palle Dahlstedt. 2017. Physical interactions with digital strings—A hybrid approach to a digital keyboard instrument. In *Proceedings of the 2017 New Interfaces for Musical Expression Conference (NIME '17)*. 115–120.

[11] Jon Forsyth, Aron Glennon, and Juan P Bello. 2011. Random access remixing on the iPad. In *Proceedings of the 2011 New Interfaces for Musical Expression Conference (NIME '11)*. 487–490.

[12] Valérian Fraisse, Marcelo M Wanderley, and Catherine Guastavino. 2021. Comprehensive framework for describing interactive sound installations: Highlighting trends through a systematic review. *Multimodal Technologies and Interaction* 5, 4 (2021), 19.

[13] Gamelatron. 2008. https://www.gamelatron.com/

[14] Lisa Rachel Gold. 2005. *Music in Bali: Experiencing music, expressing culture.* Oxford University Press.

[15] Zeyu Jin, Reid Oda, Adam Finkelstein, and Rebecca Fiebrink. 2015. MalLo: A distributed, synchronized instrument for internet music performance. In *Proceedings of the 2015 New Interfaces for Musical Expression Conference (NIME '15)*.

[16] Ajay Kapur, Michael Darling, Dimitri Diakopoulos, Jim W Murphy, Jordan Hochenbaum, Owen Vallis, and Curtis Bahn. 2011. The machine orchestra: An ensemble of human laptop performers and robotic musical instruments. *Computer Music Journal* 35, 4 (2011), 49–63.

[17] Ajay Kapur, Ge Wang, Philip Davidson, and Perry R Cook. 2005. Interactive Network Performance: A dream worth dreaming? *Organised Sound* 10, 3 (2005), 209–219.

[18] Sella Kristie, Tessa Eka Darmayanti, and Sriwinarsih Maria Kirana. 2019. Makna Motif Batik Parang Sebagai Ide Dalam Perancangan Interior. *Aksen* 3, 2 (2019), 57–69.

[19] Rafael Lozano-Hemmer. 2006. Pulse Room.

[20] José Maceda. 1986. A concept of time in a music of Southeast Asia (A preliminary account). *Ethnomusicology* 30, 1 (1986), 11–53.

[21] Margaret Mead. 1970. Children and ritual in Bali. In *Traditional Balinese Culture*. Columbia University Press, 198–211.

[22] Romain Michon and Julius O Smith III. 2014. A hybrid guitar physical model controller: The BladeAxe. In *Proceedings of the 2014 International Computer Music Conference (ICMC)*.

[23] Mogees. 2014. https://web.archive.org/web/20210624082544/https://www.mogees.co.uk/

[24] Reid Oda, Adam Finkelstein, and Rebecca Fiebrink. 2013. Towards note-level prediction for networked music performance. In *Proceedings of the 2013 New Interfaces for Musical Expression Conference (NIME '13)*. 94–97.

[25] Makoto Ono, Buntarou Shizuki, and Jiro Tanaka. 2013. Touch & activate: Adding interactivity to existing objects using active acoustic sensing. In *Proceedings of the 26th annual ACM symposium on User Interface Software and Technology*. 31–40.

[26] Francois Pachet. 2003. The Continuator: Musical interaction with style. *Journal of New Music Research* 32, 3 (2003), 333–341.

[27] Laurel Pardue, Andrew Boch, Matt Boch, Christine Southworth, and Alex Rigopulos. 2011. Gamelan Elektrika: An Electronic Balinese Gamelan.. In

Proceedings of the 2011 New Interfaces for Musical Expression Conference (NIME '11). 18–23.

[28] Miller Puckette. 2011. Playing a virtual drum from a real one. *The Journal of the Acoustical Society of America* 130, 4_Supplement (2011), 2432–2432.

[29] Pierrick Uro, Florent Berthaut, Thomas Pietrzak, and Marcelo Wanderley. 2025. Decoupling physical and virtual spaces for new collaboration strategies in co-located mixed reality instruments. In *Proceedings of the 2025 New Interfaces for Musical Expression Conference (NIME '25)*.

[30] Wayne Vitale. 2002. Balinese Kebyar Music Breaks the Five-Tone Barrier: New Composition for Seven-Tone Gamelan. *Perspectives of New Music* (2002), 5–69.

[31] Robert Wessing. 2006. A community of spirits: People, ancestors, and nature spirits in Java. *Crossroads: An Interdisciplinary Journal of Southeast Asian Studies* 18, 1 (2006), 11–111.

[32] Antonius Wiriadjaja. 2013. Gamelan Sampul: Laptop sleeve gamelan. In *Proceedings of the 2013 New Interfaces for Musical Expression (NIME '13)*. 469–470.

[33] Xiao Xiao, Paula Aguilera, Jonathan Williams, and Hiroshi Ishii. 2013. MirrorFugue iii: conjuring the recorded pianist. In *CHI'13 Extended Abstracts on Human Factors in Computing Systems*. 2891–2892.

[34] Xiao Xiao, Hayoun Noh, Adrien Lefevre, Lucy Li, Holly McKee, Alaa Algargoosh, and Hiroshi Ishii. 2025. ReMirrorFugue: Examining the emotional experience of presence and (illusory) communications across time. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–26.