Building a Harmonically Ecosystemic Machine: Combining Sonic Ecosystems with Models of Contemporary Harmonic Language

Michael Musick **Jonathan P. Forsyth** **Rachel Bittner**
Music and Audio Research Lab (MARL)
New York University
New York City, NY, USA
{ musick, jpf211, rachel.bittner }@nyu.edu

ABSTRACT

Interactive music systems are an ideal exploration ground for the testing and incorporation of decision-making algorithms and music informatics techniques into real-time applications. The interactive music system, 'The Harmonically Ecosystemic Machine; Sonic Space No. 7'*, is the compositional result of such a collaboration between music technology researchers focused on harmonic accompaniment-generating finite state transducers and sonic ecosystems as interactive music systems. This installation-based system interprets music that occurs in the physical space of a room, from a harmonic perspective, played by both the system itself and any human performers. The resulting data is used by the system to drive the ongoing, generative, interactive composition. The technical details of the system, including an overview of and motivation for the use of finite state transducers (FSTs), the process of integrating FSTs into real-time interactive music systems, the interaction paradigms established, and the experience design for the composition are presented.*

1. INTRODUCTION

The development of a system capable of generating rich musical interactions has been a goal of researchers, composers, and performers for over 40 years, dating back to the early work of composers such as Joel Chadabe and George Lewis in the 1970's [1, 2]. In general, these *interactive music systems* attempted to model various aspects of music, such as melody, harmony, rhythm, and timbre by analyzing the musical signal of an actively engaged human performer [3]. Robert Rowe identified a number of disciplines that are key to the development of interactive music systems which exhibit satisfactory levels of 'musicianship' [4, 5]; he identified the disciplines of artificial intelligence and machine listening as being particularly important.

More recently, the emergence of music information retrieval (MIR) as a distinct field within music technology has led to the development of new approaches to analyzing and modeling music. Tasks in MIR range from melody and chord transcription to genre classification and music recommendation [6]. As in the case of interactive music systems, research in MIR borrows techniques from many other fields, such as signal processing, music cognition, machine learning, and artificial intelligence.

Research in interactive music systems can make great use of the advances in music analysis and modeling forged by MIR. These new tools are important in accomplishing the ongoing efforts to create interactive music systems that are more engaging collaborators for their human participants. The level of musicianship exhibited by these systems can increase through the application of state-of-the-art techniques from MIR. Similarly, the field of MIR can benefit by applying their algorithms to challenging, real-time music systems, a process that can pose numerous challenges. For example, adapting off-line techniques for real-time systems can be problematic, and requires balancing technical and aesthetic concerns.

The recognition of these complementary interests led to the collaboration discussed here between researchers from the fields of MIR, and interactive music systems. This collaboration resulted in the composition of the interactive music system, *The Harmonically Ecosystemic Machine; Sonic Space No. 7*. This system combines Michael Musick's work in *sonic ecosystems*, a subset of interactive music systems, with Jonathan Forsyth and Rachel Bittner's research in automatic accompaniment generation using finite state machines [7, 8].

In this paper, we provide the technical details of *The Harmonically Ecosystemic Machine*, and describe some of the challenges involved in composing this system, including the trade-offs involved between integrating techniques from MIR into a real-time system. In addition, we discuss how this project furthers the evolving model for combining research efforts in MIR and interactive music systems in a mutually beneficial fashion.

The remainder of this paper is organized as follows. In section 2, we discuss the conceptual framework used to organize *The Harmonically Ecosystemic Machine*, review some background information on interactive music systems, with an emphasis on sonic ecosystems, and provide an introduction to the finite state machines used in the system. In section 3, we describe the system in detail, and in section 4, we discuss various design decisions, and how they impacted the resulting aesthetic. In section 5, we discuss our conclusions and directions for future work.

2. BACKGROUND

2.1 Interactive Music Systems

Interactive music systems, as outlined by Rowe, "are those whose behavior changes in response to musical input" [4]. This definition casts a wide net, including such systems as adaptive audio effects [9], interactive composition systems [3], and score followers [10, 11]. Additionally, there are many different sub-types of improvisational systems, i.e. systems that act as autonomous or semi-autonomous improvising agents. For example, *GenJam* [12] generates stylistically "correct" jazz improvisations over a pre-defined chord progression. Alternatively, George Lewis' *Voyager* system [2] acts as an improvising partner that conforms closely to Lewis' own particular musical aesthetic and attempts to play as an equal, intelligent musical collaborator that is taught through the performers stylistic tendencies.

Unlike *Voyager*, other improvisation generation systems attempt to model musical style more broadly, often through a combination of on-line and off-line learning. Recent research by IRCAM's music representation team has focused on multiple representations of musical information at multiple time-scales in order to create a "multi-level memory model underlying a discovery and learning process that contributes to the emergence of a creative musical agent" [13]. This work has resulted in improvisation systems such as *OMax*, which uses multi-level memory models to create a collaborative musical partner/agent. *OMax* models sequential musical information using the *factor oracle*, a type of finite state machine [13]. Similarly, Pachet's *Continuator* attempts to model the musical phrases produced by a human improviser in order to generate improvisations in a learned style by using *variable order Markov chains* (another type of finite state machine) [14].

Systems like *OMax* or the *Continuator* typically try to follow stylistic conventions of harmonic music, as has been established by 20th century practices in western music, and jazz improvisation. In addition, these systems are designed so they can perform multiple roles within the collaboration; as an accompanist to the human performer, as an equal collaborative voice, and at times as the lead of the improvisation. Since these systems are capable of playing these three roles, and because they can actively learn throughout the entire performance, the human performer is compelled to acknowledge the system as an *interactive* collaborator, and not simply as a *reactive* follower.

2.2 Sonic Ecosystems

Sonic ecosystems are an area of music making and research that Michael Musick has been pursuing with *The Sonic Spaces Project* [15–17]. These systems are considered as a distinct subset of interactive performance systems. The sonic ecosystems by Musick draw inspiration from the work of composers such as Alvin Lucier and Agostino Di Scipio; in particular Di Scipio's *Audible Eco-Systemic Interface (AESI) Project* [18], which has also influenced a number of other composers and researchers [19–21].

As with the more traditional improvising-partner style of interactive music systems mentioned above, one of the goals of a sonic ecosystem is to create a truly interactive system in which all agents have an equal opportunity to contribute musical material. However, sonic ecosystems distinguish themselves from the above interactive music systems in a number of ways. The first distinction is that sonic ecosystems typically consist of more than a single collaborative partner. Instead, these system have many "agents", whose "composed interactions" [18] are designed in a way so as to work together as a whole, much like principles of energy transfer or the so-called 'circle of life' at work in Earth's ecosystems. This results in the most important quality of these systems; that they establish an ecosystemic relationship between all of the components and agents, whereby, changes in any one of these will result in cascading changes throughout the entire system.

In order for the system to achieve a fully ecosystemic relationship, the agents need to be capable of interacting directly with each other. That is, the system must go beyond 'human-agent to digital-agent' or 'human-agent to human-agent' interaction, but also allow for 'digital-agent to digital-agent' interactions through the interface as well. Thus, musical information from digital-agents needs to feed back into the system, providing the opportunity for other digital agents to make further decisions and responses. One way this can be accomplished is by forcing all interaction between agents to occur through the same interface. In the case of sonic ecosystems, this is a sound-interface where microphones placed around the installation space capture information from both digital-agents and human-agents. The resulting relationship is one in which every agent in the system has an equal opportunity to affect the future state of every other agent in the system, human or digital.

2.3 Finite State Machines

In this section, a general overview of finite state machines is provided, followed by a brief description of their usage in speech recognition. This overview only lightly touches on the aspects of finite state machines that are most relevant to this work. For a formal detailed description of finite state machines, we refer the reader to comprehensive review articles [22, 23].

Finite state machines are graphical models which are typically represented as directed graphs consisting of nodes (states), and edges (transitions). *Finite state automata* (FSAs) are a type of finite state machine in which each edge has a single transition label. An FSA is defined by a finite *alphabet* Σ, a set of *initial states* and *final states*, and a set of *transi-*

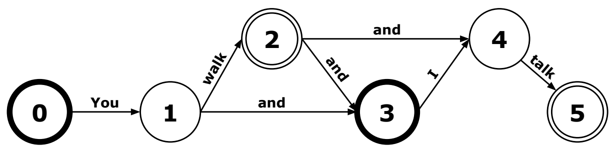

Figure 1: A simple FSA representing a set of grammatical sentences containing the words "You", "walk", "and", "I", "talk".

tions between states. The alphabet defines the set of possible values a transition label can have, and the inital and final states determine the starting and ending states of the machine. The transitions determine what states can be accessed from a given state, and in the case of a *weighted* FSA, each transition has an associated weight, which can be thought of as the probability that a transition occurs. The combination of the alphabet, initial and final states, and transitions defines the set of possible *paths* accepted by an FSA; every valid path begins at an initial state, makes transitions based on the graph structure, and terminates at a final state.

For example, Figure 1 depicts an unweighted FSA defining a simple language model. Its initial states are depicted as bold circles (states 0 and 3), and its final states are depicted as double circles (states 2 and 5). The alphabet for this FSA is $\Sigma = \{$You, walk, and, I, talk$\}$. *Paths* that would be accepted this FSA are: "You walk and I talk" (states 0, 1, 2, 3, 4, 5), "You walk and talk" (states 0, 1, 2, 4, 5), "You and I talk" (states 0, 1, 3, 4, 5), "You walk" (states 0, 1, 2), and "I talk" (states 3, 4, 5).

Finite state transducers (FSTs) are similar to FSAs, except that their edges are labeled with an input/output pair. They are defined by finite input and output alphabets, a set of initial and final states, and a set of transitions between states. Like FSAs, FSTs can be weighted. For a given *input string*, a FST will produce a set of possible *output strings*. A set of FSTs and FSAs each defining individual components of a larger system can be combined to create a model of the system.

As discussed above, FSTs and FSAs have been used to generate harmonic accompaniment to melody. The general approach is to use a FST to map melody notes to individual chords and a FSA to model chord sequences. The FST and FSA are trained separately on a dataset of melodies and harmonies, for example in "lead sheet" format. Once trained, the FST and FSA are combined into a single machine that can be used to generate a chord sequence given a new melody. In order to generate rhythmic accompaniment, we adopt a similar approach, using an alphabet to represent rhythmic patterns instead of notes or chords. A more detailed description of our approach to generating harmonic and rhythmic accompaniment can be found in "Improving and Adapting Finite State Transducer Methods for Musical Accompaniment," which serves as a companion paper to this one [8].

2.4 Organizational Principles

Rowe outlines many of the sub-disciplines that are necessary for the successful composition of interactive music systems that display high levels of musicianship [5]. He divides interactive music systems into three main components: sensing, processing, and response [4]. In this model, the sensing module is tasked with extracting information about the performance. The processing module analyzes this data and controls the state changes of the system so that it may continue collaborating with the human performer. Lastly, the response module receives control data from the processing module and handles the creation of the musical material heard during the performance.

An organizational scheme for interactive systems similar to the one just described is the "interpretative model", proposed by Blackwell and Young [24] and furthered by Bown [25]. Blackwell explains that for a system to be interactive, for example a conversation between two people (A and B), the output data from each agent should somehow depend on the received data from the other. B "must attach meaning to ([or] interpret) the input" [24] information received from A before responding with further information, and vice versa. He goes on to break down the internal process for each agent into three generally distinct modules;

- P - Listening / Feature Extractor and interpretor
- f - Reflecting / Decision Maker
- Q - Responding / Player

In the *PfQ* model, P, and Q are the interpretative functions, since they convert the input signal to an internal representation that f can process, and likewise, create output signals based on the intermediate state returned from f.

As with the earlier compositions from *The Sonic Spaces Project*, the *PfQ* organizational scheme was chosen for *The Harmonically Ecosystemic Machine*. Not only does it naturally lend itself to the types of processes that occur in an interactive music system, but it established a clear framework in which to easily test, replace, and incorporate work from the FST research project. The use of this organizational foundation for collaborations between developers is further discussed in [25].

3. THE HARMONICALLY ECOSYSTEMIC MACHINE; SONIC SPACE NO. 7

3.1 Overview

As mentioned above, the *PfQ* model allowed for easy integration of the two projects, and collaboration towards the composition. This relationship is outlined in Figure 2. In the case of composing *The Harmonically Ecosystemic Machine; Sonic Space No. 7*, the team defined the type of interactions to be composed, and the format of the data that would be passed between the modules.

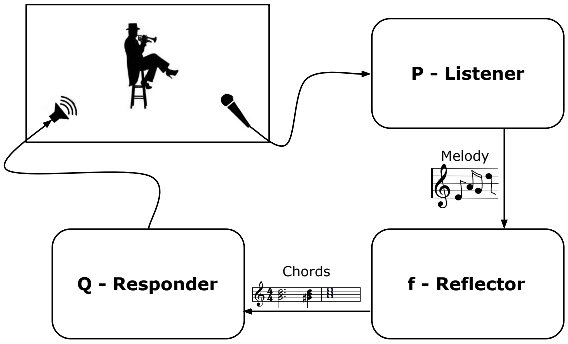

Figure 2: Top-level system diagram.

By defining the intended interactions that were to be composed, the team was able to identify a data interface for communicating between these modules. Specifically, the team defined the types of symbolic musical representation that needed to be extracted from the audio signal or created by the Decision Maker. In order to pass this data between modules, the team then only had to define an interface and the corresponding internal data format to be used between the modules and their respective software environments. This was especially useful considering the FST Project has been developed in Python, whereas the rest of the system was composed using SuperCollider. This modularized approach is crucial in facilitating this type of collaboration.

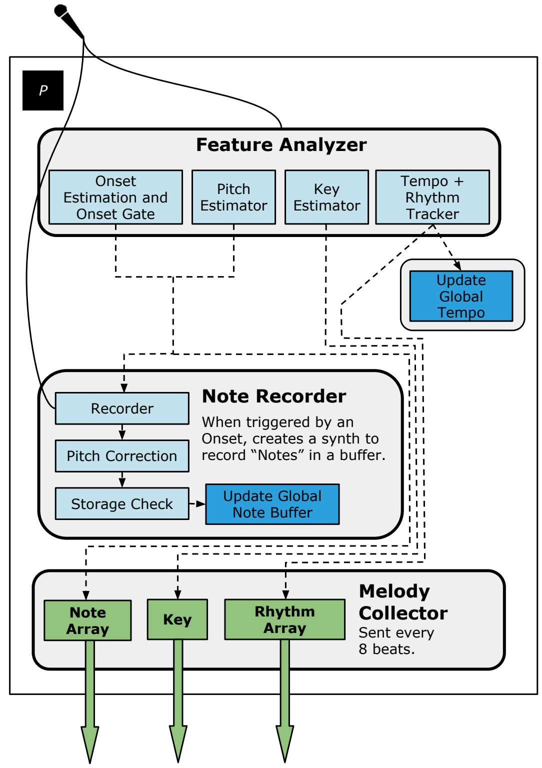

Figure 3: Overview of the *P* module.

3.2 Listening - *(P)*

The Listener (*P*) module attempts to parse emerging melodies and rhythms, extracting features and subsequently converting this information to a discrete symbolic representation. These sequences of symbolic data are then passed to the *f* module. Figure 3 represents a high level overview of the *P* module.

The "Feature Analyzer" component of the module extracts features from the incoming audio signals. The feature extractors include a noise-robust onset detector (used to determine the boundaries of note events), key and pitch estimators, and a rhythm and tempo tracker. Detected note events are collected in the "Melody Collector" as separate arrays of pitches and rhythmic values. These arrays are then sent to the Python-based *f* module every 8 beats, along with an estimate of the current key. Simultaneously, note events trigger the "Note Recorder", which captures pitch corrected notes from the audio signal and places them into an audio buffer. This audio buffer is used in the *Q* module. Finally, the global tempo is continuously updated with median-filtered values from the raw tempo estimator and is used to manage the rates of data flow between all modules.

3.3 Reflecting - *(f)*

As shown in Figure 4, the *f* module consists of two parallel sub-modules: one to generate harmonic accompaniment, and a second to generate rhythmic accompaniment. At the core of each sub-module is an FST constructed using the general framework described in section 2.3.

At the most general level, this module takes a melody array, containing 8 beats worth of individual MIDI-notes as sent from the *P* module. This array is then passed through a pre-trained FST that maps the pitch sequence to chord sequences. The symbolic chords are mapped to notes based on a predefined set of possible chord voicings. This mapping is then converted from a symbolic format to a MIDI-note sequence which is passed to the *Q* module. Similarly, the rhythmic sequence array sent from the *P* module is passed to the rhythmic accompaniment sub-module, producing a table of acceptable accompaniment rhythm sequences. Finally, both the chord and rhythm sequences are sent to the *Q* module.

FST pre-training requires a set of training data consisting of mappings from melodies to chords, and from melodic rhythms to chord rhythms. There are many possible datasets that can be used for training the harmony and rhythm generating FSTs, and the choice of dataset influences many system design decisions. We chose to train the FST's harmonic generation sub-module with the Rock Corpus Dataset [26], which provides symbolic melody labels, chord labels, beat positions, and key information for 200 rock, pop, and R&B songs. In the case of the rhythmic accompaniment generating FST, we use the rhythms from a corpus of Mozart works included in the music21 Python toolkit [27] as training data.

The FST that maps melody notes to chords takes an input melody sequence and produces a set of possible sequences

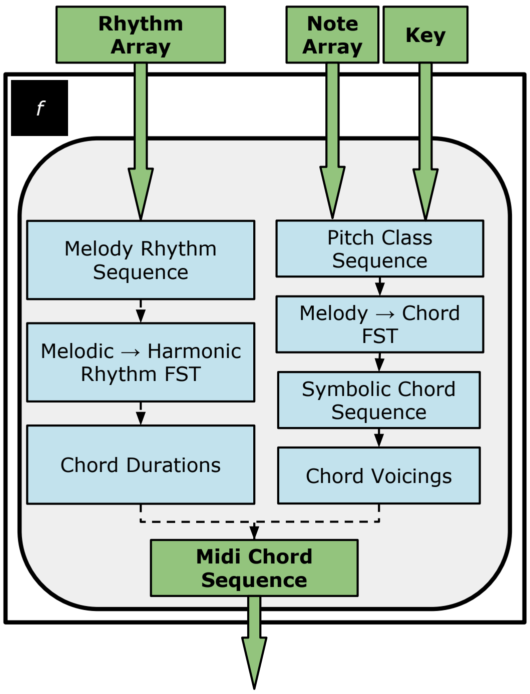

Figure 4: Overview of the f module.

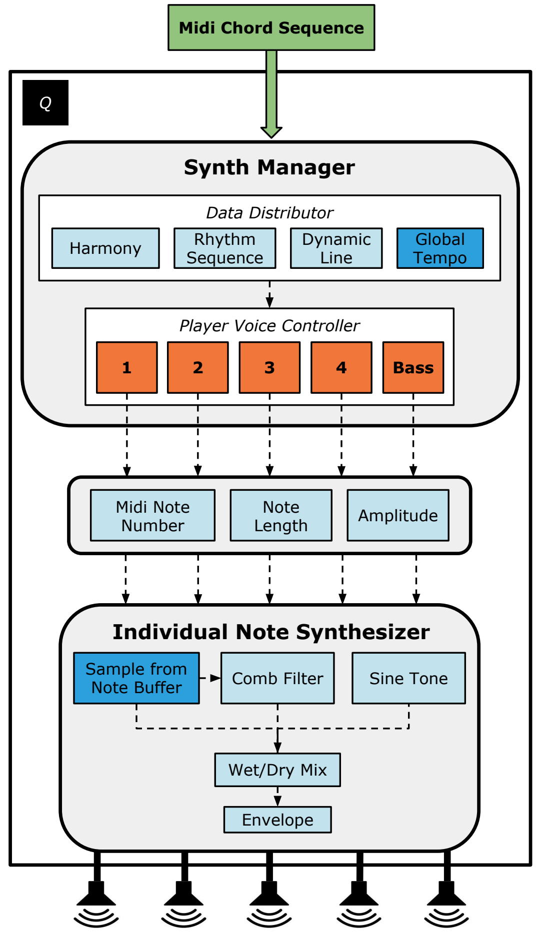

Figure 5: Overview of the Q module.

of matching chords based on the training data. Similarly, the rhythm FST produces a set of possible rhythmic sequences matching the original melodic rhythm sequence based on the training data. Note that each of these FSTs produces multiple possible output sequences for a given input – this is consistent with music, where more than one harmonization is often appropriate for a given melody – thus the output sequence is chosen randomly from the set of possible sequences.

3.4 Responding - *(Q)*

The Responder (Q) module interprets the harmonic and rhythmic sequences passed from the f module and makes musical decisions as to what notes and rhythmic sequence should be played by each of the five player agents. An overview of the Q module is shown in Figure 5. The first submodule, the "Synth Manager", splits the MIDI chord sequence across the five voices (each corresponding to a unique loudspeaker), distributing information about the pitch, duration, and volume. The information for each of the five players is sent to an "Individual Note Synthesizer", which combines a pre-recorded sample at the specified pitch from the P module's note buffer, a comb-filtered version of that audio sample, and a sine-tone at the specified pitch. This combined note is sent as a single audio signal to the corresponding player's speaker.

3.5 Room Design

The setup of the installation follows the principles established in *The Sonic Spaces Project*. The room is set up with speakers positioned in a regularly spaced circle, at 0 degrees elevation with respect to average head height (approximately 5'8"). Microphones are placed within the space, in a way that will allow them to capture multiple agents or more than just a single musical source. To encourage human participants to refrain from making music directly into the microphones, these are often placed well above head height or are disguised in sculptures. In addition, there may be "dummy" microphones placed throughout the space.

4. EXPERIENCE DISCUSSION

4.1 Interaction Goals

The overarching goal for this system, or any interactive music system that is built as an installation for people to come in and freely participate, is to create an engaging experience that encourages musicking in any way the participant feels comfortable. Within a sonic ecosystem, this is accomplished by creating interrelationships between the digital and human agents such that all of the system components are interreliant. This is also accomplished by varying the complexity of these relationships, which allows for different types of engagement with the system.

In order to further the engagement of participants, some of these relationships should be audibly apparent. Simple interaction mappings or agent relationships can serve to engage participants with limited experience in interactive music systems. Although seemingly straightforward, this is one of the compositional problems that requires artistry. If the composed relationships are too simple, as can often be the case of one-to-one mappings between input and output, participants can become bored. Likewise, there cannot be too many complex relationships, as in the case of one-to-many, or many-to-many mappings. Some variety of relationships and mappings is necessary in order to create a rich system, capable of sustaining prolonged exploration and engagement. This balance is what makes it possible for a participant to start 'knowing' the work, finding ways of continually engaging with it, and ultimately connecting with the music.

The composition of *The Harmonically Ecosystemic Machine* attempts to address these issues in a number of ways. First, the use of musical events that were recorded as audio samples from the room allows the participants to hear audible connections to the recent past. This also allows participants to observe timbral changes in the composition as the audio samples in the Note Buffer are replaced with newer musical material.

Second, the various interactions occur on different time scales. The Global Tempo and Note Buffer are both updated in the P module with minimal latency as events occur. On the other hand, the harmonic and rhythmic sequences provided by the f module are updated every eight beats, subject to the Global Tempo. Thus, the harmonic relationships between the present and the recent past remain clear. In other words, this delay time is long enough to provide meaningful data to the f module while remaining short enough that the returned harmonic sequences will be familiar and relevant to the current musical context. More importantly, the delay time allows the f module to work in near real-time, while still receiving enough data so as to make meaningful harmonic decisions. This also results in a smoother movement of harmonic sequences over a longer period of time (i.e. 16-24 beats). Various note lengths were explored for this data interface. However, eight beats proved to be the best compromise length between reducing latency in the system and providing enough melodic material to the Harmonic Decision Maker in the f module. This value is set as a global variable, and is adjusted during the tuning process of the installation system in a new space in order to verify the appropriate length for the environment.

There are numerous aspects of the f module design that can have a significant impact on the aesthetics of the final system, and thus on the engagement level of participants. Perhaps most important is the choice of training data. Since one of the stated goals is to compose an accessible and engaging system, the generated chord sequences should be interesting, yet at the same time familiar. After some experimentation, the Rock Corpus dataset was found to meet these needs. Although *The Harmonically Ecosystemic Machine* may perform in a musically unfamiliar way to some participants, the use of the familiar chord sequences typically created by the Rock Corpus dataset offers these participants something 'known' to engage with. The familiar harmonic motion that this tends to produce creates a musical environment that is familiar, yet new to participants of a wide background. The Mozart corpus was chosen as the source for the rhythmic sequence training data for the same reasons.

Finally, in order to give human participants additional ways of engaging with the ongoing composition, obvious and interesting ways of contributing music need to be present. To achieve this non-verbal invitation, a variety of simple instruments are placed throughout the room. This visual cue encourages participants to enter the space and begin exploring an instrument, thus discovering the interactions and relationships afforded by the system.

4.2 Participant Experience

During the presentations of *The Harmonically Ecosystemic Machine* that have occurred thus far, we have made a number of general observations about the participants' experiences, which were collected through conversational interviews. First, the most frequent comment received about the system was how 'fun' it is to participate with and be present in. Upon pushing for a more descriptive account, some participants mentioned that the use of recorded note samples from the room itself afforded them the opportunity to make music with re-interpreted versions of earlier moments. They also appreciated the sensitivity of the system in updating the global tempo, harmonic area, rhythmic structures, and timbral properties in relation to the collaborative music creation occurring between themselves and the digital agents. Participants also felt comfortable contributing to the system in part due to the presence of the instruments placed around the room, and in part because of the familiar sounds that the digital agents used as source material.

The system tends to associate harmonic regions with unique timbral qualities from the recent past. A small number of participants discovered this property, and as such, found that by pushing the harmonic region of the music towards a different region, they were able to shift the timbral quality of the composition.

On the other end of the response spectrum, some participants did not feel as though the system was *reactive* enough to their contributions. This can be an inherent quality of sonic ecosystems that rely on a shared acoustic interface for all human and digital agents. Essentially, sonic ecosystems provide the opportunity for all agents to equally affect change in the other agents. Unfortunately, at times this can result in a perceived sluggishness of the system in responding to sudden changes by human agents. The benefit of this interface however is that the ecosystem can continue generating musical material for a period of time without requiring additional input from human agents. This is a trade-off inherent in the composition of these systems, and the final levels of latency and sensitivity must be weighed against these qualities and desired participant experience.

Finally, participants mentioned their enjoyment of the immersive qualities of the system. Because each agent is routed to a single speaker, participants found it easier to localize the source of individual digital agents. As a result, participants have a sense of being inside a musical collaboration, with the digital agents acting as fellow collaborators.

4.3 Researchers' Experiences

This collaboration between researchers afforded a unique exploration of new interactions within sonic ecosystems. In its fifth year now as an ongoing project, *The Sonic Spaces Project* has focused on aspects of music systems that do not rely upon or relate to harmony. The primary compositional focus of prior sonic ecosystems within this project has been on the timbre qualities of music and the means of energy exchange between agents in the sonic interface. The opportunity to use a pre-trained FST in the f module, while maintaining the previous aesthetic principles of the project was important in exploring previous assumptions about participant experience and interaction with these systems.

One question of interest at the start of this collaboration was how the inclusion of harmonic musical properties would change the focus and experience of participants. Although a formal study has yet to be completed, it would seem from the presentations that have occurred thus far that adding a harmonic aspect to one of the interaction layers of the system did change the focus and experience of participants considerably. Rather than focusing on how individual sounds morph and may relate between agents over time, participants were more concerned with the harmonies being produced by the system. More specifically, participants were concerned with whether the system was harmonically following or *reacting* to what they personally were doing. The addition of harmonic relationships was a detriment to the timbral links to the past created by the global note sample array. This layer tended to be missed by participants unless highlighted by a member of the research team.

This collaboration created an opportunity to explore ways of incorporating the fundamental principles of *The Sonic Spaces Project*, as outlined in [17], to a harmonically focused system. Also, by specifying the desired interactions of the system with the other collaborators, as well as the data transmission protocol between modules, clear objectives were defined for the composition of the P and Q modules, while allowing the f module to optionally be treated as a 'black box'.

Incorporating the FST-based harmonic accompaniment generation system into a real-time context was a goal from the start of the accompaniment project. The original inspiration for the development of the accompaniment generation system was in the practice of *live-looping*. Thus, the composition of *The Harmonically Ecosystemic Machine*, a real-time interactive music system with an emphasis on the use of re-sampled and re-contextualized audio, provided an ideal opportunity to explore adapting the accompaniment system to a real-time context.

This process of adaptation forced the accompaniment system researchers to reconsider various technical aspects of the system. For example, the real-time nature of *The Harmonically Ecosystemic Machine* required that the harmonic accompaniment generation process was reasonably efficient. This requirement imposed certain constraints on the system, such as the size of the FST, which in turn imposed limits on various system parameters and design possibilities. The process of incorporating the accompaniment project into a larger system raised other practical concerns as well, such as the ease of deployment and effective communication strategies with the other modules of the larger system. Just as specifying the data transmission protocols afforded the P and Q modules freedom to develop independently, this same specification allowed for the f module to be developed knowing it would easily interface with the larger system. Finally, this process also encouraged the researchers to develop new techniques to generate rhythmic accompaniment.

5. CONCLUSION

This collaboration created many opportunities for the individual researchers to explore various questions related to sonic ecosystems, finite state machines, artificial intelligence, interactive music systems, and experience design. Perhaps more importantly, it served as another step in establishing practices for the mutually beneficial collaboration between the research disciplines of MIR and interactive music systems. These disciplines have natural connections, shared problems, and the demonstrated ability to support each others' continued development.

This paper has outlined a number of the more pertinent technical and aesthetic details about *The Harmonically Ecosystemic Machine; Sonic Space No. 7.* [1] As a unique iteration in both of these research projects, this composition served as a vehicle to work out new problems, present the research projects in new ways, and receive feedback from a multidisciplinary group. The final system itself is a valuable composition, which, as evidenced from the installation presentations thus far, offers an engaging experience, on multiple levels, for participants willing to explore the interconnected relationships created by this type of sonic ecosystem.

6. REFERENCES

[1] J. Chadabe, "Interactive Composing: An Overview," *Computer Music Journal*, vol. 8, no. 1, pp. 22–27, 1984. [Online]. Available: http://www.jstor.org/stable/3679894

[2] G. E. Lewis, "Too Many Notes: Computers, Complexity and Culture in Voyager," *Leonardo Music Journal*, vol. 10, pp. 33–39, Dec. 2000. [Online]. Available: http://www.mitpressjournals.org/doi/abs/10.1162/096112100570585#.U0xdz8a9NFw

[3] J. Chadabe, "Some Reflections on the Nature of the Landscape within which Computer Music Systems are Designed," *Computer Music Journal*, vol. 1,

[1] For more information about *The Harmonically Ecosystemic Machine*, including recorded audio examples, the score, and code, please visit `http://steinhardt.nyu.edu/marl/research/sonic_spaces`.

no. 3, pp. 5–11, 1977. [Online]. Available: http://www.jstor.org/stable/10.2307/3679605

[4] R. Rowe, *Interactive Music Systems: Machine Listening and Composing*. The MIT Press, 1992.

[5] ——, *Machine Musicianship*. The MIT Press, 2004.

[6] A. Lerch, *An Introduction to Audio Content Analysis: Applications in Signal Processing and Music Informatics*. John Wiley & Sons, 2012, vol. 5.

[7] J. P. Forsyth and J. P. Bello, "Generating Musical Accompaniment Using Finite State Transducers," in *Proc. of the 16th Int. Conf. on Digital Audio Effects (DAFx). Maynooth, Ireland (Sept 2013)*, 2013.

[8] J. P. Forsyth, R. Bittner, M. Musick, and J. P. Bello, "Improving and Adapting Finite State Transducer Methods for Musical Accompaniment," in *Proceedings of the International Computer Music Conference (ICMC)*, 2015.

[9] V. Verfaille, U. Zolzer, and D. Arfib, "Adaptive Digital Audio Effects (A-DAFx): A New Class of Sound Transformations," *Audio, Speech, and Language Processing, IEEE Transactions on*, vol. 14, no. 5, pp. 1817–1831, 2006.

[10] C. Raphael, "Music Plus One and Machine Learning," in *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, 2010, pp. 21–28.

[11] A. Cont, "ANTESCOFO: Anticipatory Synchronization and Control of Interactive Parameters in Computer Music." in *International Computer Music Conference (ICMC)*, 2008, pp. 33–40.

[12] J. A. Biles, "Evolutionary Computation for Musical Tasks," in *Evolutionary computer music*. Springer, 2007, pp. 28–51.

[13] G. Assayag, "Creative Symbolic Interaction: Keynote talk," in *40th International Computer Music Conference (ICMC) and 11th Sound and Music Computing (SMC) conference*, no. September, Athens, Greece, 2014, pp. 1–6.

[14] F. Pachet, "The Continuator: Musical Interaction with Style," *Journal of New Music Research*, vol. 32, no. 3, pp. 333–341, 2003.

[15] M. Musick, "Sonic Spaces Project," Masters Thesis, University of Michigan, 2012. [Online]. Available: http://michaelmusick.com/2012/08/02/thesis-paper/

[16] ——, "Examining the Analysis of Dynamical Sonic Ecosystems: in Light of a Criterion for Evaluating Theories," in *40th International Computer Music Conference (ICMC) and 11th Sound and Music Computing (SMC) conference*, Athens, Greece, 2014, pp. 154–161. [Online]. Available: http://quod.lib.umich.edu/i/icmc/bbp2372.2014.026/1

[17] M. Musick and T. H. Park, "Timbral Hauntings: An Interactive System Re-Interpreting the Present in Echoes of the Past," in *40th International Computer Music Conference (ICMC) and 11th Sound and Music Computing (SMC) conference*, Athens, Greece, 2014, pp. 254–261. [Online]. Available: http://quod.lib.umich.edu/i/icmc/bbp2372.2014.040/1

[18] A. Di Scipio, "Sound is the Interface: From Interactive to Ecosystemic Signal Processing," *Organised Sound*, vol. 8, no. 03, pp. 269–277, Apr. 2003. [Online]. Available: http://www.journals.cambridge.org/abstract_S1355771803000244

[19] O. T. Green, "User Serviceable Parts: Practice, Technology, Sociality and Method in Live Electronic Musicking," Ph.D Thesis, Electroacoustic Music, City University London, 2013.

[20] O. Green, "Audible Ecosystemics as Artefactual Assemblages: Thoughts on Making and Knowing Prompted by Practical Investigation of Di Scipio's Work," *Contemporary Music Review*, vol. 33, no. 1, pp. 59–70, May 2014. [Online]. Available: http://www.tandfonline.com/doi/abs/10.1080/07494467.2014.906698

[21] S. Waters, "Performance Ecosystems: Ecological Approaches to Musical Interaction." in *EMS: Electroacoustic Music Studies Network*, De Montfort/Leicester, 2007, pp. 1–20. [Online]. Available: http://webuser.fh-furtwangen.de/~friedm/PerformEcosystems.pdf

[22] E. Vidal, F. Thollard, C. De La Higuera, F. Casacuberta, and R. C. Carrasco, "Probabilistic Finite-State Machines – Part I," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 27, no. 7, pp. 1013–1025, 2005.

[23] ——, "Probabilistic Finite-State Machines – Part II," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 27, no. 7, pp. 1026–1039, 2005.

[24] T. Blackwell and M. Young, "Self-organised music," *Organised Sound*, vol. 9, no. 02, pp. 123–136, 2004.

[25] O. Bown, "Experiments in Modular Design for the Creative Composition of Live Algorithms," *Computer Music Journal*, vol. 35, no. 3, pp. 73–85, Sep. 2011. [Online]. Available: http://www.mitpressjournals.org/doi/abs/10.1162/COMJ_a_00070

[26] T. De Clercq and D. Temperley, "A Corpus Analysis of Rock Harmony," *Popular Music*, vol. 30, no. 01, pp. 47–70, 2011.

[27] M. S. Cuthbert and C. Ariza, "music21: A Toolkit for Computer-Aided Musicology and Symbolic Music Data," in *Proceedings of the International Symposium on Music Information Retrieval*, vol. 11, 2010, pp. 637–42.