

EXTENDING THE MUSICAL EXPERIENCE

FROM THE PHYSICAL TO THE DIGITAL AND BACK

Gil Weinberg

1. Introduction

It is widely perceived that the computer has enriched and advanced the art form of music. Digital technology brought new palettes of sounds, composition techniques, and production methods; innovations in digital compression and distribution changed music consumption and listening practices; for performers, novel musical instruments and controllers have been developed based on a variety of sensing, interaction, and mapping approaches. But after more than two decades of research in computer music, a fundamental question must be asked – has digital technology truly innovated and enriched the expressive, emotional, and creative core of the musical experience? It is not clear that the answer to this question is as positive as we music technologists would like to think.

During the last ten years, inspired and motivated by the prospect of innovating the core of the musical experience, I have explored a number of research directions in which meaningful use of digital technology bears the promise of revolutionising the medium. The research directions identified – gestural expression, collaborative networks, and constructionist learning – can lead to musical experiences that cannot be facilitated by traditional means. The first direction builds on the notion that through novel sensing and mapping techniques, new expressive musical gestures can be discovered that are not supported by current acoustic instruments. Such gestures, unconstrained by the physical limitation of acoustic sound production, can provide infinite possibilities for expressive and creative musical experiences for novice as well as trained musicians. The second research direction utilises the digital network in an effort to create new collaborative experiences, allowing players to take an active role in determining and influencing not only their own musical output but also that of their co-performers. By using the network to interdependently share and control musical materials in a group, musicians can combine their musical ideas into a constantly evolving collaborative musical activity that is novel and inspiring. The third research direction utilises constructionist learning, which bears the promise of revolutionising music education by providing hands-on access to programmable music making. Through interaction with physical computational objects, learners can construct personally meaningful musical artifacts that enhance and deepen their learning.

While facilitating novel musical experiences that cannot be achieved by traditional means, the digital nature of these research directions often leads to flat and inanimate speaker-generated sound, hampering the physical richness and visual expression of acoustic music. In my most recent work, therefore, I attempt to combine the benefits of digital computation and acoustic richness, by exploring the concept of “robotic musicianship”. I define this concept as a combination of musical, perceptual, and social skills with the

capacity to produce rich acoustic responses in a physical and visual manner. The robotic musicianship project aims to combine human creativity, emotion, and aesthetic judgment with computational capabilities, allowing human and robotic players to cooperate and build off one another's ideas. A perceptual and improvisatory robot can best facilitate such interactions by bringing the computer into the physical world both acoustically and visually.

In this paper I will describe my projects portraying a musical journey that was initiated by my interest in extending acoustic music with digital technology and reached its most recent period by investigating the enhancement of digital music through physical-acoustical means. Each station in this journey presents a different set of novel expressive and creative possibilities along with a set of limitations and constraints imposed by technology.

2. Related Work, Goals, and Challenges

The field of New Interfaces for Musical Expression¹ has received significant interest in recent years as researchers and musicians explore new sensing techniques, design approaches, mapping schemes, and sound generation methods to enhance and enrich musical expression. Research in this area can be categorised into two main areas – Imitated and Augmented Instruments, and Alternate Controllers. Building on the vast repertoire of familiar musical gestures, researchers have created imitated and augmented versions of traditional instruments such as percussions, strings and woodwinds, among others. Alternative ways to play music have also been explored by using various sensing and mapping techniques such as in non-contact instruments wearable music and alternate tangible controller. Most of these instruments, however, have been created for particular compositions (usually by the inventor) and have been effective only within specific aesthetics boundaries. Only few controllers have shown durability and adaptability to multiple compositions in a variety of musical styles. Inspired by the tradition of great versatile acoustic instrument such as the piano, one of the main goals of my work was to develop controllers that are durable, versatile, and adaptable to multiple compositions, styles, and playing techniques.

The second area of related work is in the field of Interconnected Musical Networks (IMNs) – live performance systems that allow players to influence, share, and shape each others' music in real-time. Such systems, whether they operate in one physical space or over a wide-area network, provide an interdependent framework that can lead to rich social and musical experiences that are not supported by traditional group play. The development of

1 <<http://www.nime.org>>

IMNs since the 1950s has been connected to the development of technological innovations – from John Cage’s early experimentations with interconnected transistor radios through the use of networked PCs by groups like the League of Automatic Music Composers and the Hub, to the current proliferation in collaborative Internet music. These experiments, however, usually require advanced musical skills and understanding by players and audiences, and often lead to inaccessible “high art” musical outcome. More recent collaborative musical installations for novices on the other hand, tend to simplify the musical experience for novices and are not geared to interdependently connect between novices and professionals. To address this gap, my work attempts to explore novel interdependent musical interactions that would provide both novices and experts with rich and inspiring, yet intuitive and easy to follow, collaborative musical experiences.

The educational goal of my research is informed by related work in the field of constructionist learning. The constructionist approach emphasises the unique ability of digital technology to provide personal and configurable learning experiences to a wide variety of learners. The approach was conceived by Seymour Papert, who demonstrated how learning is most effective when students construct personally meaningful technological artifacts. Other researchers have elaborated on Papert’s ideas, showing how interaction with digital physical objects enhances children’s and adults’ learning. In music, however, little has been done to develop constructionist systems that attempt to connect between figural expressive musical experiences and formal aspects of theory and technique. In conventional music education systems, when music students are introduced to formal theory, certain important expressive aspects that came naturally in the early figural mode are temporarily hidden when learners try to superimpose analytical knowledge upon felt intuitions. My work attempts to utilise constructionist-learning methods to bridge the gap between the figural and formal learning modes through hands-on interaction with programmable musical controllers.

And lastly, I introduce the concept of *robotic musicianship*, taking up Rowe’s concept of *machine musicianship*. In this research area, scholars develop interactive systems that analyse, perform, and compose music with computers based on theoretical foundations in fields such as music theory, computer music, music cognition, and artificial intelligence. Several effective approaches for the design of such interactive musical systems have been explored over the years by researchers and musicians such as Dannenberg², Cope³, Lewis⁴, Pachet⁵, and others. Such digital interactive systems, however,

2 Dannenberg 1984

3 Cope 1996

4 Lewis 2000

5 Pachet 2002

are limited by the inanimate and flat nature of their digital musical reproduction. Current research directions in musical robotics, on the other hand, focus mostly on sound production and rarely address social aspects such as listening, analysis, group improvisation, or collaboration. Both “robotic instruments”⁶ – mechanically automated devices that can be played by live musicians or triggered by pre-recorded sequences – and “anthropomorphic robots”⁷ – hominoid robots that attempt to imitate the action of human musicians – function mostly as mechanical apparatuses that follow deterministic rules. The motivation for establishing the field of *robotic musicianship* is to develop robots that can produce rich acoustic sound and visual cues, while utilising computational power and techniques of machine musicianship that are not possible with traditional acoustic instruments.

3. The Projects

3.1 The Musical Playpen (1997-1998)

The Musical Playpen was the framework for my preliminary experimentation with gestural musical interaction in a constructionist-learning environment. The instrument was designed for toddlers and infants in an effort to explore whether very young children can participate in a meaningful, active



Fig. 1. A child playing in the Musical Playpen

musical experience. The environment allows young children to control two high-level musical aspects – contour and rhythmic stability – in an environment which is both familiar and fun: a 1.5-x-1.5-m playpen filled with 400 colourful plastic balls (Fig. 1). The playpen was designed to generate musical responses in correlation to children’s activity. Players’ movements around the playpen propagated from ball to ball and triggered four piezo-electric

sensors that were hidden inside four balls, one in each corner of the playpen. The balls’ ability to transmit hits to neighboring balls, combined with the sensors’ high sensitivity allowed for almost any delicate movement around the playpen to be captured by at least one sensor. The analog signal was then digitised and sent to a Macintosh computer running Max/MSP where it was mapped to musical output played from speakers below the playpen. Two opposite corners were mapped to control the melodic contour of an Indian raga,

6 For example, see Dannenberg et al. 2005, Jordà 2002, Singer et al. 2004.

7 For example, see Takanishi et al. 1998, Toyota 2004.

so that the more energetic the players' movements in these corners were, the higher the played Indian raga pitches became. Children could therefore create melodic phrases and manipulate their curves by changing the intensity of their body movements in these corners. Player's physical activity in the other two corners were mapped to an algorithm that controlled the tempo, rhythmic variation, and timbre of percussive sequences in an effort to provide access to controlling rhythmic stability. The more energetic the players were near these corners, the more versatile and uneven the rhythmic values became. The tempo curve also fluctuated more sharply, as did the rate of timbral change.

A number of observation sessions were conducted with the playpen at MIT and at the Boston Children's Museum from 1998 to 1999. These sessions have shown a wide range of responses to the environment and the high-level musical control that it offered. For example, a 1-year-old infant started her session by triggering a sequence of notes as she was placed near one of the melodic curve corners. The infant looked in the direction of the sound source and tried to move her hand towards that corner, seemingly trying to repeat the music she heard. When she succeeded and another melodic phrase was played, she smiled, took one ball and tried to shake it, obviously without audible results. Frustrated, she then threw the ball towards a rhythmic corner, generating a short percussive sequence. She approached this corner while moving her torso back and forth, laughing when discovering that her movements controlled the music. After a short break the infant started to move her body again back and forth, gradually accelerating her movements, generating less and less stable percussive sequences. Only after repeating this behaviour in another corner did the infant seem to be ready to use more expressive, less restricted gestures all over the playpen.

These responses can indicate that with the right instruments and controls, young children can have access to spontaneous, expressive music-making as well as to more serious and thoughtful musical explorations. These findings encouraged me to develop a new set of instruments, which I entitled "The Squeezables", in an effort to continue and develop models for high-level musical control, and to explore novel methods for networked group collaboration with older players, who can express and discuss their impression of the experience.

3.2 The Squeezables (1998-1999)

In the Squeezables project, I attempted to add the concept of musical networks to my initial interest in gestural controllers and constructionist education. The goal of the project was to allow a group of players, novices and proficient musicians, to interdependently collaborate in constructing a

meaningful musical composition using unconventional expressive gestures. The instrument consisted of six squeezable and retractable gel balls mounted on a small podium, which players could simultaneously squeeze and pull to manipulate a set of low- and high-level musical percepts. The combination of pulling and squeezing allowed players to utilise familiar and expressive gestures and to control multiple synchronous and continuous musical parameters. Several materials were tested and for the final prototype, soft gel balls were chosen, which proved to be robust and responsive, providing a sense of force feedback control that derived from the elastic qualities of the gel. Buried inside each ball was a 0.5-x-2.0-cm plastic block covered with five pressure sensors, protected from the gel by an elastic membrane. The analog pressure values from these sensors were transmitted to a digitiser and converted to MIDI. Pulling gestures were sensed by six variable resistors installed under the table. An elastic band connected to each ball added opposing force to the pulling gesture, helping to retract the balls back onto the tabletop (Fig. 2).



Fig. 2. Three networked players play *The Squeezables*

In an effort to evaluate the high-level algorithms in the instrument, a number of straightforward mappings were designed to control relatively low-level musical parameters. For example, one of the balls formed a one-to-one connection between squeezing and pulling gestures to the modulation rate and range of two low-frequency oscillators, respectively. For other balls higher-level algorithms were developed to control percepts such as contour and stability. For example, pulling and squeezing gestures of the “Arpeggiator” ball controlled a combination of musical parameters including tempo, pitch commonality, dissonance and rhythmic variation, so that the more the ball was squeezed and pulled, the more unstable an arpeggiated sequence became. To facilitate a coherent hierarchical interconnected interaction, the balls were divided into five accompaniment balls and one melody soloist. The five accompaniment balls provided players with autonomous control – no input from the other balls influenced their output. However, these balls’ output was mapped not only to the accompaniment parameters but also to transform the sound of the “melody” ball. While pulling the “melody” ball manipulated its own contour so that the higher it was pulled, the higher the melodic curve became. The actual pitches, as well as the MIDI velocity, duration and pan values, were determined by the level of pulling and squeezing of the accompaniment balls. This allowed the accompaniment balls to “shape” the character of the melody while maintaining a comprehensive scheme of interaction among themselves.

To experiment with these mappings I composed a short piece for three players. The piece, which was featured in *Ars Electronica 2000*⁸, starts with a high-level of instability and builds gradually towards a repetitive rhythmic peak. Special notation was created for the piece – two continuous graphs were assigned to each one of the six balls. One graph indicated the level of squeezing over time and the other indicated the level of pulling. The process of writing and performing the piece served as a useful tool for evaluating the mapping and sensing techniques used. In addition, discussions were held with novices and professionals who played the instrument. In general, children and novices were more inclined to prefer playing the balls that provided high-level control such as contour and stability. They often stated that these balls allowed them to be more expressive and less analytical. Proficient musicians, on the other hand, often found the high-level control somewhat frustrating, because it did not provide them with direct and precise access to specific desired parameters. Some experts complained that their personal interpretation of the high-level controllers for stability differed from the one implemented in designing the instrument. Both novices and professional players found the multiple-channel synchronous control expressive and challenging and the pulling and squeezing gestures comfortable and intuitive.

These gestures allowed delicate and easily learned control of many simultaneous parameters, which was especially compelling for children and novices. The organic and responsive nature of the balls was one of the features mentioned as contributing to this expressive experience. When asked about the interdependent networked connections, one melody ball player described her experience as a constant state of trying to expect the unexpected. To another player, the experience felt like controlling an entity with a life of its own. In a manner similar to chamber music group interaction, body and facial gestures served an important role in coordinating the accompaniment players' gestures and establishing an effective outcome. Such collaborations turned out to be especially compelling for children, who found the accompaniment balls conducive to social interaction, intuitive and easy to play with. Some complaints were made, however, regarding the difficulty for individual accompaniment players to create their own musical phrases without being constantly subjected to interdependent transformation from the group. Other criticism addressed the lack of discrete input, which prevented players from generating and controlling specific musical events in detail.

8 <<http://www.aec.at/festival2000/>>

3.3 The Musical Fireflies (1999-2000)

The Musical Firefly project was designed to address some of the weaknesses in the Squeezables. In particular, it aimed to facilitate a more discrete and autonomous interaction that would allow for clearer interaction schemes and more focused constructionist-learning goals. The project attempted to provide players with expressive hands-on experiences that can be easily transformed into an analytical and formal exploration of music and mathematics. Through simple tapping gestures players could input rhythmic patterns and embellish them in real-time by adding multiple rhythmic layers. This functionality provided players with figural and formal familiarisation with musical concepts such as accents, beats, patterns, and timbre. During the multi-player interaction, a wireless network was formed between Fireflies, which allowed players to synchronise patterns and trade instrument sounds. This interactive group experience was designed to lead to deeper internalisation of advanced musical concepts such as the correlation between monorhythmic and polyrhythmic structures. Access to and manipulation of LOGO code for customising the controllers provided an introduction to MIDI programming and electronic sound. Advanced players could, therefore, deepen their learning experience by reprogramming the controllers and adjusting their functionality to match personal musical interests and abilities.

The 3D printed Musical Firefly's case was designed to be held by two hands while thumb-tapping two top-mounted buttons. Signals from the buttons were sent to an embedded "Cricket" Microchip PIC microprocessor. An infrared communication port allowed for communication with other Fireflies as well as for downloading LOGO based application programs. The




Fig. 3. Two players interact with each other with the Musical Fireflies

played rhythmic patterns were converted into musical messages using Cricket LOGO general MIDI commands and sent through the Cricket's serial bus port to the MidiBoat – a small General Midi circuit that supported up to 16 polyphonic channels, 128 melodic timbres and 128 percussive timbres. The audio from the MidiBoat was then sent to the top-mounted speaker.

Interaction with the Musical Fireflies occurred in two distinct and sequential modes – the Single Player Mode, where players converted numerical patterns into rhythmical structures, and the Multi Player Mode, where collaboration with other players enhanced the basic rhythmic structures into polyrhythmic compositions (Fig. 3). In Single Player

Mode, players could trigger and play with two default percussive sounds. The left button triggered accented notes and the right button triggered non-accented notes. The patterns of accented and non-accented notes were recorded and after two seconds of inactivity, played back in a loop, using an adjustable default tempo. This activity provided players with a tangible manner of entering and listening to the rhythmical output of any numerical pattern they envisioned, leading to an immediate conceptualisation of the mathematical-rhythmical correlation. For example, Figure 4 depicts the playing of the numerical pattern 4 3 5 2 2:



Fig. 4. A pattern of accented and non-accented notes as played by the Musical Fireflies. ● = Accented note played by the left button; ○ = non accented note played by the right button

During playback, players could enter a second layer of accented and non-accented notes in real-time, using a different timbre. Each tap on a button triggered a note aloud and recorded its quantised position so that the pattern became part of the rhythmic loop. Pressing both buttons simultaneously at any point stopped the playback and allowed the player to enter a different pattern. In Multi Player Mode, when two loop playing Fireflies “saw” each other (i.e., when their infrared signals were exchanged), they automatically synchronised their rhythmic patterns. (A similar interaction occurs when the Firefly insects synchronise their light pulses to communicate in the dark). This activity provided participants with a richer, more complex rhythmical composition and allowed for an interactive introduction to polyrhythm. Figure 5 depicts how a 7 beat pattern played by one Firefly and a 4 beat pattern played by another diverge and converge as the patterns go in and out of phase every 28 beats, the smallest common denominator:




Fig. 5. Two patterns (7/4 and 4/4) played by two Fireflies divergence and convergence as they go in and out of phase every 28 beats

While the two Fireflies were synchronised, players could also initiate a “Timbre Trade” in which instrument sounds were exchanged between the devices. Pressing either the left or right button traded both layers of the accented or non-accented timbres respectively. Each Firefly continued to play its original pattern using the new received timbre. This interaction provided players with a higher-level of musical abstraction as they separated the rhythmical aspect of the beat from the timbre in which it was played. Because the Fireflies network became richer after the interaction (i.e., each instrument contained four different timbres) the system encouraged collaborative play where players were motivated by trading, collecting and playing games by sending and receiving different timbres from their peers.

Observations of play sessions with the Musical Fireflies have been conducted followed by discussions with the players. Participants were asked about the expressive and the educational aspects of the session as well as for their suggestions for improvements. A software version of the application was prepared and tested. Both novices and experienced users found the concrete aspects of playing with a physical object compelling in comparison with the graphical user interface of the software version, mentioning the unmediated connection that was formed with the instrument as contributing to the creation of personal connection with their music they created. Listening to the music from distinct physical sources also helped players to follow the interaction in a more coherent manner in comparison to listening to computer speakers. The observations and interviews also led to the identification of points for improvement and future work. For example, it was clear that the focus on a specific constructionist learning activity hampered the open-ended expressive gestural interaction goal of the project. Moreover, the simple interaction using only two discrete buttons and the low-quality MIDI sounds led to a disappointing musical outcome, consisting mostly of monotonous interlocking clicks with no pitch, time-based rhythmic values, rests, or continuous transformation. The network interaction in multi-user mode, while effective for learning, did not provide a satisfactory collaborative experience. The restricted interconnectivity of the system, where discrete timbre-trading was the only interpersonal act, did not provide long-lasting rich play value and led players to lose interest in the interaction after a few trades. In addition, due to the limitations imposed by the line-of-sight infrared communication, the application only allowed for synchronisation and timbre trading between two players at a time. Many interviewees expressed their wishes to interact and collaborate in larger groups comprised of several simultaneous players.

3.4 The Beatbugs / “Nerve” (2001-2003)

For the Beatbug project, new hardware and software applications were developed in an effort to address the weaknesses identified in the Musical Fireflies. The binary buttons were replaced with a piezo electric sensor that could sense hit strength, providing more expressive physical interaction through large full-arm drumming gestures. The single user application was enhanced to record rhythmic values, rests, pitches, and amplitudes, allowing for more versatile and expressive musical input. Two new bend sensors were added to the design, allowing players to continuously modify and transform the recorded musical phrases using low- and high-level transformation algorithms (Fig. 6). In addition, the embedded MIDIBoat was replaced with a high-quality software synthesiser, which significantly enhanced sound quality and versatility. Several important enhancements were also made to improve the multi-user collaborative interaction. The network was enhanced to support up to eight simultaneous Beatbugs, while coloured LEDs were installed in each Beatbug to help convey complex multi-user interactions in a visual manner. The interpersonal application was improved to provide longer lasting collaborative interactions, allowing players to continuously develop each other’s music by bending and manipulating the Beatbug antennae. In order to support these improvements, the new Beatbugs communicated with each other through wires via a central computer system, which was titled the “Nerve Center”. To showcase the improved system, a musical composition was composed, titled “Nerve”, which was presented in workshops and concerts as part of the Tod Machover’s Toy Symphony project.

In an effort to provide a familiar and fun interface for children and novices, the “Nerve” Beatbug was designed as a bug, having a speaker for a mouth, two bend-sensors for antennae, and a velocity-sensitive piezoelectric sensor on its back. White and coloured LEDs mounted in its translucent shell provided visual feedback when hit or played through. An embedded Microchip PIC microcontroller was responsible for reading input from the sensors, controlling the LEDs, and communicating with the central system via tail-like cable that carried MIDI, trigger, audio, and power. The piezo electric sensor measured when and how hard it was hit, while the two antennae allowed for subtle control over different aspects of the sound. Bending the antennae caused a proportional change in the colour of three LED clusters, and a ring of white LEDs flashed each time the bug was hit, providing additional visual feedback



Fig. 6. Manipulating the Beatbugs antennae

to the player and audience. The embedded processor was responsible for operating the sensors and LEDs, while the central computer system controlled the actual musical interactions and behaviours. The “brain” of the system was written in Max/MSP environment. Controlling all of the behaviour from the central computer made it easy to quickly experiment with a broad range of interaction schemes. Similarly, sound synthesis occurring on the central computer and played through the corresponding Beatbug’s speaker, provided high quality sound with an embedded, self-contained feel. For the software synthesiser, ‘Reason’ by Propellerhead was chosen, providing a broad palette of timbres and continuous control over multiple sound parameters. Up to eight Beatbugs could be connected to one central rack, which consisted mostly of standard off-the-shelf equipment including an audio interface, a MIDI interfaces, an 8-channel amplifier, and a mixer. The only non-standard device in the system was a custom patch box, which provided power to the bugs and converted the 10-pin connector in each cable to MIDI in, MIDI out, trigger, and audio in (Fig. 7).




Fig. 7. The Nerve Beatbug system's schematics

Similarly to the Musical Fireflies, players interacted with the “Nerve” Beatbug in two distinct modes – Single Player Mode, and Multi Player Mode. In Single Player Mode, each player could enter a short rhythmic pattern over a predefined metronome beat. The system automatically played back the recorded pattern in a loop through the corresponding Beatbug’s speaker. A quantisation algorithm pushed the notes towards the closest quarter, eighth or triplet note. While the entered pattern was playing back, the player could manipulate the pattern by bending the two antennae. The left antenna continuously transformed the pitch and timbre using a variety of predefined scales and audio effects. The right antenna added rhythmic ornamentation to the pattern by controlling the values, length, accentuation, and feedback level of a delay line. The goal of these transformation algorithms was to allow players to modify the pattern but to keep the feel of the original motif, supporting the “motif-and-variation” nature of the interaction. In Multi Player Mode players could form large-scale collaborative compositions by interdependently sharing and continuously developing each other’s motifs. Each Beatbug player could play a rhythmic motif that was then automatically sent through the stochastic computerised “Nerve Center” to another player in the group. The receiving player could decide whether to further develop the received motif (by continuously manipulating pitch, timbre, and rhythmic elements with the two bend sensor antennae) or to keep the motif in his or hers personal bug (by entering and sending a newly generated motifs to a different random player in the group). The antennae transformations were recorded and layered in each cycle until a new pattern was entered. The tension between the system’s stochastic routing scheme and the players’ improvised real-time decisions led to an interdependent, dynamic, and constantly evolving musical outcome. In a different section of Multi Player Mode, after all players entered their patterns, the system awaited a series of simultaneous hits by all players that led to random segmentation of the participants to sub-groups, allowing players to interdependently collaborate with a gradually growing number of co-players.⁹



Fig. 8. A Beatbug workshop at MIT Media Lab

During 2002-2003 the “Nerve” Beatbugs were featured in workshops and concerts in Berlin, Dublin, Glasgow, Boston and New York in collaboration with local symphonies and educational programs (Fig. 8). During each week-long workshop, children and orchestra members were introduced to

⁹ See a video clip of the interaction as performed in concert at <<http://www.cc.gatech.edu/~gilwein/videos/Glasgow%20-%20Concert.mov>>.

the Beatbugs, explored the system, and rehearsed towards a public concert. The workshops also featured a new constructionist pedagogy developed in collaboration with Kevin Jennings. The pedagogy was designed to allow players to physically create and phrase rhythmic patterns and transform them by employing melodic, timbral, and rhythmic contours. The balance among aural, kinesthetic and social modalities provided the children with a rich and highly immersive musical environment. A report by Project Zero from Harvard's Education School said that "[the project] provided an overwhelmingly positive experience either from the musical, social and personal standpoint... the experience provided a good foundation on which to build one's musicianship, social skills, self-confidence, and general learning dispositions focusing, listening, and practicing."

Several problems and areas for improvement became apparent as well. The musical mappings in Single Player Mode, although more versatile than in the Musical Fireflies, were still limited and unsatisfactory for many proficient musicians, who expressed their interest in creating and manipulating more advanced and non-quantised melodic and harmonic musical content. Novices too showed interest in controlling more sophisticated musical material even if they could not create it themselves. In multiplayer interactions,



Fig. 9. Large play gestures in a "Nerve" concert, Cambridge, MA

the velocity sensing piezoelectric sensor and the large scale of the system encouraged players to use wide playing gestures and expressively point to indicate their actions to each other and to the audiences (Fig. 9). However, while these large gestures brought elements of visual expression and excitement to the performance,

they were not sensed by the central system and therefore did not have audible consequences. In terms of hardware, it was clear that the central system was too large and complex, and that the 18-unit rack was not easily portable. An additional hardware weakness was the durability of the bend sensor antennae, which proved to be fragile, especially when large groups of energetic children experimented with the system during week-long workshops.

3.5 *iltur* (2003-2005)

The *iltur* project utilised an improved version of the Beatbug controllers, which were enhanced both in hardware and software in an effort to address the weaknesses observed in Nerve. Hardware improvements included replacing the unreliable bend sensors with robust Hall effect sensors, installing 2D accelerometers to sense larger and more expressive arm gestures, and reducing the size and complexity of the system. The software was rewritten to address users' requests to control and manipulate advanced melodic and harmonic content in a more expressive and gestural manner. The new application supported interaction between Beatbug players and proficient musicians, allowing Beatbug players to record live input from MIDI and acoustic instruments and to respond by transforming the recorded material gesturally, creating motif-and-variation call-and-response routines on the fly. The central computer host was programmed to analyse MIDI and audio signals and to allow Beatbug players to personalise the analysed material using a variety of transformation algorithms. Capturing and personalising richer musical content through expressive gestures gave Beatbug players the opportunity to create a more sophisticated musical outcome, while forming elaborate musical dialogs with their peers.

The main hardware improvement in the *iltur* Beatbugs was the addition of the 2D accelerometers. The accelerometers were used to sense tilting and shaking gestures, providing the central system with information regarding players' large arm movements. Hardware improvements were also made in an effort to make the antennae more robust, utilising Hall effect sensors and magnets mounted under the antennae. This electromagnetic sensing method proved to be robust and effective, although it provided lower bending resolution in comparison with the original resistance-based bend sensors. Other hardware improvements addressed the system size and portability. As opposed to the complex 18-unit rack Nerve system, the new *iltur* system, utilised a laptop instead of a desktop, a software mixer instead of a physical one, and no MIDI drum controller, as audio from the piezoelectric sensors was captured directly through an audio interface. The system, therefore, was housed in a small 6-unit rack (Fig. 10).

Play gestures and interaction in *iltur* were modified to allow for recording, triggering, and manipulation of MIDI and audio in real-time. Recording was conducted by simultaneously bending both antennae while tapping the Beatbug. The system then segmented the recorded phrases, looking for sections of silence in the MIDI and/or audio buffers. The audio Beatbugs were programmed to detect onset notes, pitches, and amplitudes in real-time. The analysis algorithm was optimised for brass instruments and was used successfully with instruments such as trumpet, trombone and saxophone. Onset




Fig. 10. The iltur Beatbug system's schematics

identification and segmentation of MIDI was trivial due to the discrete nature of the MIDI protocol. After the system recorded and segmented the captured musical input, players could immediately trigger the recorded phrase by tapping the Beatbug again. Hit velocities were mapped to different segments in the phrase, allowing players to rearrange the recorded motifs. Two synthesis methods – Wavetable Synthesis and Granular Synthesis – were used for re-triggering audio. The Wavetable technique provided close resemblance to the sound of original recording but suffered from noise artifacts during continuous transformations. Granular Synthesis, on the other hand, provided harsher sounds in comparison to the original recording but allowed for smoother continuous transformation. A number of different mapping schemes were experimented with for antennae bending and accelerometer-based gestures. Some of these algorithms utilised direct mappings between continuous gestures and fundamental musical aspects such as pitch, volume and tempo. Other mapping approaches allowed for the manipulation of higher-level musical percepts such as melodic similarity or rhythmic density. Shaking gestures were most successful when mapped to control vibrato and tremolo effects, while antennae manipulations were effective in controlling pitch. When interacting with a MIDI instrument, Beatbug players could also trigger the recorded

motif in inversion and retrograde by tapping the Beatbug while bending the left or right antennae, respectively. The audio Beatbugs allowed players to control transformations such as pitch bending, speed alteration, and filtration, through a combination of bending, tilting, and hitting gestures. During group interaction, players could trade their motifs by simultaneously hitting the Beatbug while bending one of the antennae. Receiving players could then further transform the phrase and send it back to their peers. In comparison to the random involuntarily routing scheme in *Nerve*, *iltur* players could trade their motifs only when simultaneously agreeing to synchronise their gestures. Three Jazz compositions were written for the *iltur* system and performed in cities such as Atlanta, San Diego, Miami, Vancouver, and Jerusalem. *iltur 1* featured MIDI interaction, *iltur 2* focused on audio transformation and manipulation, and *iltur 3* introduced group interaction and motif trading. Voice manipulation experimentations were also conducted, allowing Beatbug players to interact with a hip-hop vocalist.¹⁰



Fig. 11. *iltur 3* audio Beatbug players interact with a brass section (left) and a hip-hop vocalist (right) in Jerusalem, Israel

Observations of and discussion with *iltur* players led to a number of findings regarding the improved Beatbug functionalities. For example, it was clear the *iltur* Beatbugs were more effective than the *Nerve* Beatbugs in providing richer musical experiences for individuals through a larger set of expressive gestures and more complex melodic and harmonic transformations. The new application also led to more meaningful and versatile collaborations between novices and professional musicians. Both players and audiences perceived the new accelerometer-based gestures as intuitive, expressive, and visually compelling. However, the introduction of gesture combinations (such as hitting the Beatbug while bending the antenna) was problematic for novices and children, who found it physically and mentally challenging. Novices and children also found the higher-level transformation algorithms (such as musical density and stability) less intuitive to control and preferred the simple and predictable one-to-one mappings between gestures and low-level musical

¹⁰ See videos at <<http://www.cc.gatech.edu/~gilwein/iltur.htm>>.

aspects. More proficient musicians, on the other hand, preferred to interact with the high-level musical operations, stating that these encouraged them to concentrate on the correlation between their actions and the musical output.



Fig. 12. Interaction between two iltur 3 MIDI Beatbug players

In general, the effectiveness of the experience was closely related to the musical and harmonic context of the compositions. Due to segmentation and audio stretching, in a harmonically structured composition it was difficult for players to improvise while following the harmonic progression. Many players, therefore, preferred free musical structures,

stating that open-ended experience posed less boundaries and allowed more creativity and expression.

3.6 Haile (2004-2007)

The instruments and controllers discussed above explored different ways in which meaningful embodiment of technology can enhance the musical experience by facilitating new expressive gestures, networked group collaborations and constructionist learning. Although these projects provided satisfying results, the instruments were limited by the electronic reproduction and amplification of sound through speakers, which did capture the richness of acoustic sound. My most recent project – an interactive robotic percussionist named Haile – addressed this limitation by utilising a mechanical apparatus that converts digital musical instructions into acoustic and physical generation of sound. Haile was developed in an effort to bring together the advantages of computational power with the expression and richness of creating acoustic sound using physical and visual gestures.

The project aimed to combine that are not possible by humans with rich sound and visual gestures that cannot be reproduced by speakers in an effort to facilitate new musical experiences, and new music, that cannot be conceived by acoustic or means.

As part of the project, a robotic percussionist that listened to and analysed live musical input in real-time and reacted by generating relevant, but

at times surprising, acoustic responses was developed. The project posed challenges in areas such as perception modeling, mechanics, and interaction design. In perception, the main challenge was to implement models for low- and high-level musical percepts, allowing the robot to develop a meaningful representation of the music it listened to. In mechanics the challenge was to develop a dexterous robotic apparatus that would translate perceptually based performance algorithms into a rich acoustic and visually informative performance. In interaction design, our aim was to develop performance algorithms that would enable the robot to collaborate with human players in a meaningful and intuitive manner, using transformative and generative methods both sequentially and synchronously.

In order to support familiar interactions with human players, Haile's design is anthropomorphic, utilising two percussive arms that can move to different locations and strike with varying velocities (Fig. 13). The first prototype was designed to play a Native American Pow Wow drum – a multi player instrument that supported the collaborative nature of the project. For pitch-oriented applications, the robot was later adjusted to play a one-octave xylophone. In order to match the aesthetics of these musical instruments, Haile was constructed from wood using a CnC cutting machine. Metal joints were designed to allow shoulder and elbow movement as well as leg adjustability for different instrument heights. While attempting to create an organic look for the robot, it was also important that the technology was not completely hidden, so that co-players could see and understand the robot's operation. The mechanical apparatus was therefore left uncovered and LEDs were embedded on Haile's body, providing an additional representation of the mechanical actions. Haile's right arm was designed to play fast notes, while the left arm was designed to produce larger and more visible motions that produce louder sounds. Both arms could adjust the strikes sound in two manners: different pitches were achieved by striking the instruments in different locations, and volume was adjusted by hitting with varying velocities. To move to different vertical positions, each arm employed a linear slide, a belt, a pulley system, and a potentiometer to provide feedback. Unlike robotic drumming systems that allow hits at only a few discrete locations, Haile's arms moved continuously over a distance of 10



Fig. 13. Haile, the perceptual robotic percussionist, listens to and interacts with a human player

inches (movement timing is 250 ms. from end to end). The right arm's striking mechanism was loosely based on a piano hammer action and consisted of a solenoid driven device and a return spring. The right arm stroked at a maximum speed of 15 Hz, faster than the left arm's maximum speed of 11 Hz. However, the right arm did generate a wide dynamic range or provided easily noticeable visual cues, which limited Haile's expression and interaction potential. The left arm was designed to address these shortcomings, using larger visual movements, and a more powerful and sophisticated hitting mechanism.

The first phase of the project aimed at facilitating rhythmic collaboration between human drummers and Haile, addressing aspects such as rhythmic perception, improvisation, and interaction design. Perceptual models were developed for low- and high-level rhythmic percepts, from beat and density analysis, to rhythmic stability and similarity perception. Some relatively low-level perceptual modules included beat analysis, where domain detection was followed by autocorrelation of tempo and phase, and density analysis, where we looked at the number of note onsets per time unit to represent the density of the rhythmic structure. Higher-level rhythmic analysis modules were also developed for percepts such as rhythmic stability, based on research by Desain, et al.¹¹, and rhythmic similarity based on Tanguiane's survey¹². The stability model calculated the relationship between pairs of adjacent note durations, rated according to their perceptual expectancy based on three main criteria: perfect integer relationships were favoured, ratios had inherent expectancies (i.e., 1:2 was favoured to 1:3 and 3:1 was favoured to 1:3), and durations of 0.6 seconds were preferred. The similarity rating was derived from Tanguiane's binary representation, where two rhythms are first quantised, and then given a score based on the number of note onset overlaps and near-overlaps.

The main challenge in designing the rhythmic interaction with Haile was to implement the perceptual modules in a manner that would lead to an inspiring human-machine collaboration. The approach taken was based on a theory of interdependent group interaction in interconnected musical networks. At the core of this theory is a categorisation of collaborative musical interactions in networks of artificial and live musicians based on sequential and synchronous operations with centralised and decentralised control schemes. Based on this framework, six interaction modes were developed: Imitation, Stochastic Transformation, Perceptual Transformation, Beat Detection, Simple Accompaniment, and Perceptual Accompaniment. These interaction modes utilised different perceptual modules and were embedded

11 Desain et al. 2002

12 Tanguiane 1993

in different combinations in interactive compositions and educational activities. In the first mode, Imitation, Haile merely repeated what it heard based on its low-level onset, pitch, and amplitude perception modules. Players could play a rhythm and after a couple of seconds of inactivity Haile imitated it in a sequential call-and-response manner. Haile used one of the arms to play lower pitches close to the drumhead centre and the other arm to play higher pitches close to the rim. In the second mode, Stochastic Transformation, Haile improvised in a call-and-response manner based on players' input. Here, the robot stochastically divided, multiplied, or skipped certain beats in the input rhythm, creating variations of users' rhythmic motifs while keeping their original feel. Different transformation coefficients were adjusted manually or automatically to control the level of similarity between humans' motifs and Haile's responses. In the Perceptual Transformation mode, Haile analysed the stability level of users' rhythms, and responded by choosing and playing other rhythms that had similar levels of stability to the original input. In this mode Haile automatically responded after a specified phrase length. Imitation, Stochastic Transformation, and Perceptual Transformation were all sequential interaction modes that formed decentralised call-and-response routines between human players and the robot. Beat Detection and Simple Accompaniment modes, on the other hand, allowed synchronous interaction where humans played simultaneously with Haile. In Beat Detection mode, Haile tracked the tempo and beat of the input rhythm using complex domain detection function and autocorrelation, which led to continuously refined assumptions of tempo and phase. A simpler, yet effective, synchronous interaction mode was Simple Accompaniment, where Haile played pre-recorded MIDI files so that players could interact with it by playing their own rhythms or by modifying elements such as drumhead pressure to modulate and transform Haile's timbres in real-time. This synchronous centralised mode allowed composers to feature their structured compositions in a manner that was not susceptible to algorithmic transformation or significant user input. The Simple Accompaniment mode was also useful for sections of synchronised unisons where human players and Haile played together. Perhaps the most advanced mode of interaction was the Perceptual Accompaniment mode, which combined synchronous, sequential, centralised and decentralised operations. Here, Haile played simultaneously with human players while listening to and analysing their input. It then created local call-and-response interactions with different players, based on its perceptual analysis. In this mode amplitude and density perceptual modules were utilised – while Haile played short looped sequences (captured during the Imitation and Stochastic Transformation modes) it also listened to and analysed the amplitude and density curves of human playing. It then modified its looped sequence, based on the amplitude and density coefficients of the human players. When the

rhythmic input from human players was dense, Haile played sparsely, providing only the strong beats and allowing humans to perform denser solos. When humans played sparsely, on the other hand, Haile improvised using dense rhythms that were based on stochastic and perceptual transformations. Haile also responded in direct relationship to the amplitude of human players so that the louder humans played, the stronger Haile played to accommodate the human dynamics, and vice versa.¹³

As a creative outcome for these interactive applications, two compositions were written for the system, each utilised a different set of perceptual and interaction modules. The first composition, titled *Pow*, was written for one or two human players and a one-armed robotic percussionist. It served as test



Fig. 14. A performance of *Jam'aa* in Odense, Denmark

case for Haile's early mechanical, perceptual, and interaction modules. The second composition, titled *Jam'aa* ("gathering" in Arabic), built on the unique communal nature of the Middle Eastern percussion ensemble, attempting to enrich its improvisational nature, call-and-response routines, and virtuoso solos with algorithmic transformation and human-robotic interactions (Fig. 14). *Jam'aa*, was commissioned by Hamaabada Art Centre In

Jerusalem, and later performed in invited and juried concerts in France, Germany, Denmark, and the United States.¹⁴

As part of our effort to expand the exploration of robotic musicianship into pitch and melody, Haile was later adapted to play a pitch-based mallet instrument. A one-octave xylophone was built for this purpose to accommodate Haile's mechanical design – the left arm covered a range of 5 keys while the right arm, whose vertical range was extendable, covered a range of 7 keys. (Fig. 15). Following the idiom "listen like a human, improvise like a machine", computational models for melodic similarity were developed ("listen like a human") as the fit function of a genetic algorithm based improvisation engine ("improvise like a machine"). The algorithmic responses were based on the analysed input as well as on internalised knowledge of contextually relevant

13 See a video excerpts of some of the interaction modes at <<http://www.cc.gatech.edu/~gilwein/Haile.htm>>.

14 See a video excerpts from *Jam'aa* at <<http://coa.gatech.edu/~gil/RoboraveShort.mov>>.

material. The algorithm fragmented MIDI and audio input to short phrases. It then attempted to find a “fit” response by evolving a pre-stored human-generated population of phrases using a variety of mutation and crossover functions over a variable number of generations. At each generation, the evolved phrases were evaluated by a fitness function that measured similarity to the input phrase, and the least fit phrases in the database are replaced by members of the next generation. A unique aspect in this design was the reliance on a pre-recorded human-generated phrase set that evolved over a limited number of generations. This allowed musical elements from the original phrases to mix with elements of real-time input to create hybrid, and at times unpredictable, responses for each given input melody. Two compositions were written for the system and performed in concerts in




Fig. 15. Haile's adaptation for xylophone



Fig. 16. A performance of Svobod in Copenhagen, Denmark

Atlanta and Copenhagen. In the first piece, titled “Svobod”, a piano and a saxophone player freely improvised with a semi-autonomous robot (Fig. 16). The second piece, titled “iltur for Haile”, involved a tonal musical structure utilising genetically driven and non-genetically driven interaction schemes, as the robot performed autonomously with a jazz quartet.¹⁵

¹⁵ See a video clip of iltur for Haile at <<http://www.coa.gatech.edu/~gil/iltur4haile.mov>>.

References

- Aimi, Roberto (2002): *New Expressive Percussion Instruments*, M.S. Thesis, MIT Media Laboratory.
- Aimi, Roberto/Young, Diana (2004): »A New Beatbug: Revisions, Simplifications, and New Directions«. In: *Proceedings of the 2004 International Computer Music Conference*, San Francisco: International Computer Music Association.
- Bamberger, Jeanne (2000): *Developing Musical Intuitions*, New York: Oxford University Press.
- Barbosa, Alvaro (2003): »Displaced Soundscapes: A Survey of Network Systems for Music and Sonic Art Creation«. *Leonardo Music Journal* 13, 53-59.
- Bischoff, John/Gold, Rich/Horton, Jim (1978): »Microcomputer Network Music«. *Computer Music Journal* 2(3), 24-19.
- Buchla, Don (2005): »A history of Buchla's musical instruments«. In: Sidney Fels/Tina Blaine (Eds.), *New Interfaces for Musical Expression, NIME-05, Proceedings*, Vancouver: University of British Columbia, Media and Graphics Interdisciplinary Center (MAGIC), 1.
- Cope, David (1996): *Experiments in Music Intelligence*, Madison WI: A-R Editions.
- Cycling74 (2005): *Max/MSP* <<http://www.cycling74.com>> (last access: October 2007).
- Dannenberg, Roger B. (1985): »An On-line Algorithm for Real-Time Accompaniment«. In: William Buxton (Ed.), *Proceedings of the 1984 International Computer Music Conference*, San Francisco: International Computer Music Association, 193-198.
- Dannenberg, Roger B./Brown, Ben/Zeglin, Garth/Lupish, Ron (2005): »McBlare: A Robotic Bagpipe Player«. In: Sidney Fels/Tina Blaine (Eds.), *New Interfaces for Musical Expression, NIME-05, Proceedings*, Vancouver: University of British Columbia, Media and Graphics Interdisciplinary Center (MAGIC), 80-84.
- Desain, Peter/Honing, Henkjan (2002): »Rhythmic stability as explanation of category size«. In: *Proceedings of the International Conference on Music Perception and Cognition*, Sydney: UNSW, CD-Rom.
- Feldmeier, Mark C./Malinowski, Mateusz/Paradiso, Joseph A. (2002): »Large Group Musical Interaction Using Disposable Wireless Motion Sensors«. In: *Proceedings of the 2002 International Computer Music Conference*, San Francisco: International Computer Music Association, 83-87.
- Gardner, Howard (1983): *Frames of Mind: The Theory of Multiple Intelligences*, New York: Basic Books.
- Gresham-Lancaster, Scot (1998): »The Aesthetics and History of the Hub: The Effects of Changing Technology on Network Computer Music«. *Leonardo Music Journal* 8, 39-44.
- Jennings, Kevin (2003): »Toy Symphony: An International Music Technology Project for Children«. *Music Education International* 2, 3-21.

- Jordà, Sergi (2002): »Afasia: The Ultimate Homeric One-man multimedia-band«. In: Eoin Brazil (Ed.), *New Interfaces for Musical Expression*, NIME-02, Proceedings, Dublin: Media Lab Europe, 132-137.
- Levin, Golan (2001): »Dialtone: A Telesymphony«. Online available: <<http://www.flong.com/projects/telesymphony/>> (last access: October 2007).
- Lewis, George E. (2000): »Too Many Notes: Computers, Complexity and Culture in Voyager«. *Leonardo Music Journal* 10, 33-39.
- Machover, Tod (1992): *Hyperinstruments: A Progress Report 1987-1991*, Cambridge: MIT Media Laboratory.
- Machover, Tod (2004): »Shaping Minds Musically«. *BT Technology Journal* 22(4), 171-179.
- Martin, Fred/Mikhak, Bakhtiar/Silverman, Brian (2000): »MetaCricket: A designer's kit for making computational devices«. *IBM Systems Journal* 39, 3-4.
- Mathews, Max V. (1991): »The Radio Baton and Conductor Program, or: Pitch, the Most Important and Least Expressive Part of Music«. *Computer Music Journal* 15(4), 37-46.
- Nikitina, Svetlana (2004): »Toy Symphony Report«, Cambridge: Harvard University School of Education.
- Pachet, Francois (2002a): »The continuator: Musical interaction with style«. In: *International Computer Music Association (Ed.), Proceedings of the 2002 International Computer Music Conference*. San Francisco: International Computer Music Association.
- Papert, Seymour (1980): *Children, Computers, and Powerful Ideas*, New York: Basic Books.
- Paradiso, Joseph A. (2004): »Wearable Wireless Sensing for Interactive Media«. In: *Proceedings of the 1st International Workshop on Wearable and Implantable Body Sensor Networks*, 2004, 30-31.
- Patten, James/Recht, Benjamin/Ishii, Hiroshi (2002): »Audiopad: a tag-based interface for musical performance«. In: Eoin Brazil (Ed.), *New Interfaces for Musical Expression*, NIME-02, Proceedings, Dublin: Media Lab Europe, 11-16.
- Proceedings of International Conferences on New Interfaces for Musical Expression (NIME), <<http://www.nime.org/>> (last access: October 2007).
- Resnick, Mitchel/Martin, Fred/Sargent, Randy/Silverman, Brian (1996): »Programmable Bricks: Toys to Think With«. *IBM Systems Journal* 35(3-4), 443-452.
- Rowe, Robert (2004): *Machine Musicianship*, Cambridge: MIT Press.
- Scheirer, Eric D. (1998): »Tempo and beat analysis of acoustic musical signals«. In: *Journal of the Acoustical Society of America* 103(1), 588-601.
- Singer, Eric/Feddersen, Jeff/Redmon, Chad/Bowen, Bil (2004): »LEMUR's Musical Robots«. In: Yoichi Nagashima/Michael J. Lyons (Eds.), *New Interfaces for Musical Expression*, NIME-04, Proceedings, Hamamatsu: Shizuoka University of Art and Culture, 181-184.

- Smith, Joshua R./Strickon, Josh (1998): »The MIDIBoat«. Online available: <<http://web.media.mit.edu/~jrs/minimidi/>> (last access: October 2007).
- Steiner, Nyle A. (2004): »The Electronic Valve Instrument (EVI), an electronic musical wind controller for playing synthesizers«. *Journal of the Acoustical Society of America* 115(5), 2451-2451.
- Takanishi, Atsuo/Maeda, Maki (1998): »Development of Anthropomorphic Flutist Robot WF-3RIV«. In: *Proceedings of the 1984 International Computer Music Conference*, San Francisco: International Computer Music Association, 328-331.
- Tanguiane, Andranick S. (1993): *Artificial Perception and Music Recognition*, New York: Springer.
- Theremin, Leon (1996): »The Design of a Musical Instrument Based on Cathode Relays«. *Leonardo Music Journal* 6, 49-50.
- Toyota (2004): »Trumpet Robot«. <<http://www.toyota.co.jp/en/special/robot>> (last accessed October 2007).
- Waiswiz, M. (1985): »The Hands: A Set of Remote MIDI Controllers«. In: *Proceedings of the 1985 International Computer Music Conference*. San Francisco: International Computer Music Association, 313-318.
- Weinberg, Gil (2001): »The Squeezables: Toward an Expressive and Interdependent Multi-player Musical Instrument«. *Computer Music Journal* 25(2), 37-45.
- Weinberg, Gil (2005): »Interconnected Musical Networks – Toward a Theoretical Framework«. *Computer Music Journal* 29(2), 23-39.
- Weinberg, Gil/Lackner, Tamara M./Jay, Jason (2000): »The Musical Fireflies – Learning About Mathematical Patterns in Music Through Expression and Play«. In: *Proceedings of the 2000 Colloquium on Musical Informatics*. LAquila, Italy: Istituto GRAMMA, 146-149.
- Weinberg Gil/Aimi, Roberto/Jennings, Kevin (2002): »The Beatbug network: a rhythmic system for interdependent group collaboration«. In: Eoin Brazil (Ed.), *New Interfaces for Musical Expression, NIME-02, Proceedings*, Dublin: Media Lab Europe, 1-6.
- Weinberg, Gil/Driscoll, Scott (2005): »iltur: Connecting Novices and Experts Through Collaborative Improvisation«. In: Sidney Fels/Tina Blaine (Eds.), *New Interfaces for Musical Expression, NIME-05, Proceedings*, Vancouver: University of British Columbia, Media and Graphics Interdisciplinary Center (MAGIC), 17-22.

Acknowledgements

The projects described in this paper would not have been possible without the valuable contribution of colleagues and students at MIT Media Lab and Georgia Tech Music department. In particular, I would like to thank Seum Lim Gan, Roberto Aimi, Tamara Lackner, Jason Jay, Scott Driscoll, Travis Thatcher, Mark Godfrey, Alex Rae, and John Rhoads for their indispensable contribution for the development of the musical instruments and applications. I would also like to thank Tod Machover, director of the Hyperinstrument group at MIT Media Lab, and Frank Clark, director of the Music Department at Georgia Tech for their support.