Realtime Representation and Gestural Control of Musical Polytempi

Chris Nash and Alan Blackwell
Rainbow (Interaction & Graphics) Research Group, University of Cambridge
Computer Laboratory, 15 JJ Thomson Avenue, Cambridge, CB3 0FD, UK
+44 (0)1223 334678
{christopher.nash, alan.blackwell}@cl.cam.ac.uk

ABSTRACT
Over the last century, composers have made increasingly ambitious experiments with musical time, but have been impeded in expressing more temporally-complex musical processes by the limitations of both music notations and human performers. In this paper, we describe a computer-based notation and gestural control system for independently manipulating the tempi of musical parts within a piece, at performance time. We describe how the problem was approached, drawing upon feedback and suggestions from consultations across multiple disciplines, seeking analogous problems in other fields. Throughout, our approach is guided and, ultimately, assessed by an established professional composer, who was able to interact with a working prototype of the system.

Keywords
Tempo, polytempi, performance, composition, realtime, gesture

1. INTRODUCTION
Although intricate and complex musical processes involving rhythm, melody and harmony are to be found in most musical genres, the use of and conventions relating to tempo are less adventurous [16]. It has only been the last century or so that has seen composers, such as Steve Reich and Conlon Nancarrow, experiment with simultaneous musical parts bearing differing tempi [9]. As regards experiments in musical time, the notion of polytempi is crucially different from the relatively more common concepts of polyrhythm and polymetre, which both rely on simple integer divisions of the bars or beats in the piece. In contrast, the multiple simultaneous tempi of polytempo music leads to situations where the bar lines and beats of each part in the piece are themselves incongruent. The timing relationships between the events in each part can no longer be thought of, or expressed in, simple integer fractions (e.g. 3 in the time of 2, or 3/2 vs. 6/4), but instead become irrational.

A number of explanations can be volunteered for the paucity of polytempo use in the modern musical repertoire. For the average listener: the barrage of incongruous notes resulting from multiple, misaligned passages of music could be argued to hold limited aesthetic appeal. For the performer: such perceptions of anarchy, chaos and randomness may seem ironic, as the musician attempts to follow the composer's explicit, highly-ordered and inflexible timing directions, unable to rely on implicit or explicit timing cues from other musicians or a conductor [19]; unable to rely on a universal, steady pulse of bar or beat. And, finally: if the performer struggles to manage an individual part of the piece, then the composer's task of developing and imagining the complete, combined performance becomes an almost impossibly hard mental operation.

In conventional music, an audience is invariably attuned to global tempo variations within a piece. When introducing a simultaneous part with a differing tempo, an extra dimension is added to the audience's perception of the performance – the explicit interplay of parts in respect of time.

Figure 1. Perceived synchronisation in phase music.

Due to the periodic nature of much of the world's rhythms [3], there are various points where disjoint parts can appear more or less temporally-aligned, so that the perceived effect is determined not only by the absolute musical offset, but also relative factors. For example, in Figure 1, the parts start in-sync and gradually diverge because of differing tempi. Initially, the divergence is small enough so that the listener can still corrupt their perception of the musical events onto a single time scale, dismissing the offset as they might a digital chorus effect, acoustic echo or performance prosody [9]. After time, the offset increases and the two parts are more easily separated, becoming harder to align perceptually. Yet, by Bar 4, the *absolute* offset is approximately one beat, and thus the music can be aligned about the beat. Continued, such alignment occurs *relative* to other points in the bar, as well as divisions of the beat, and inevitably aligns relative to the bar itself.

The varying incongruity of notes can be seen to form a temporal harmony, where perceived aligned and misaligned episodes correspond to consonance and dissonance respectively. For centuries, these concepts have been powerful tools levered by composers in their engagement with tonal harmony; in the typical case, dissonance giving way to consonance, to provide resolution [14]. As such, and like dissonant harmonies, the average listener's aversion to apparently cacophonic, misaligned music serves only to reinforce the potential for temporal resolutions.

Harmony and pitch have been studied, codified and notated in a variety of ways to enable performance by musicians and experimentation by composers, yet few notations have arisen to explicitly express temporal relations [16]. In our research, we apply computer technology and interactive notations to tackle these remaining problems. Whereas the problem of directing and coordinating performers is unquestionably also a matter of notation (be it paper, digital, aural, static or interactive), this paper principally focuses on the earlier, pre-requisite stage in the process – the composer's creation of the music. For our purposes, the "super-human virtuosity" currently required of polytempi performers can be provided by the computer [4].

2. BACKGROUND

The simple concept described in the Introduction and Figure 1 underpins the phase music of Steve Reich [15]. "Piano Phase" (1967) contains two musical parts with the same melodic and rhythmic content, but slightly differing (yet constant) tempi. The parts start together, then gradually diverge, or 'phase', in musical time, producing moments of dissonance and consonance, as the parts become more or less aligned. Reich's interface to this process began with a tape machine, playing two looped tapes of the phrase at different speeds. Subsequently, and owing to the relative simplicity and repetitive nature of the musical content, he was able to carry the idea to the piano, whereupon two exceptionally disciplined and practiced performers can play the music live. With the exception of the tape speed settings, general performance directions and the looped phrase itself, however, the piece is not fully-scored, but is instead an example of the generative or procedural specification of music. Notably, it is difficult to inspect or manipulate specific, individual notes or events in the performance.

Conlon Nancarrow, a contemporary of Reich, took a different approach to the problems of notation and performance, replacing the human pianist with a pianola (or player piano), notating his music on the paper roles used by the machine [9]. Unlike score notation, the roles represent time linearly and the piano's mechanism eventually afforded the opportunity to dynamically vary tempo within a part. Unlike Reich, Nancarrow's pieces tended not to rely on the phasing of musical events in repetitive parts, but on a grander plan of having a single, climatic point of synchrony.

Alejandro Viñao [1], an established modern-day composer, has been much inspired by Nancarrow's efforts, and now brings a more personal perspective and motivation to our research, joining us as Composer in Residence at the Computer Laboratory. For more than 30 years and in a variety of areas and centres of research (including IRCAM and MIT's Media Lab), he has sought technologies to help him express his musical ideas. Yet, his appropriation of technologies and methods in conventional music practice force him to an unsatisfying compromise when it comes to exploring polytempi. Using scored music, Alejandro divides the bar into the finest performable resolution (e.g. 1/32nd notes), and uses varying note accents and stresses to give the impression of multiple tempi. Even though he admits such methods do not produce true polytempi, Alejandro manages to create pieces that are nonetheless able to present impressions of temporal harmony, with temporally dissonant passages resolving to consonance. Furthermore, his reliance on more established working practices affords him greater flexibility in instrumentation, arrangement and performance.

Both Nancarrow and Reich effectively used technology to address problems with the use of polytempi, but were both forced to pre-calculate and prescribe the tempo variations long in advance of performance; waiting hours, days or weeks to hear the result of their writing. In all three cases, the composers are forced to limit their creativity in some way, be it temporal freedom, dynamism, note-to-note control or instrumentation.

Approaches to managing complex musical timings tend to focus on performance requirements. Ghent [7] is one of the earlier attempts to use audio cues (e.g. multiple metronomes) for individual musicians. Ligeti [12] uses a similarly audio-based method. Such techniques isolate the musician from the ensemble and, more importantly, the part from the piece, which is not only incompatible with the composer's requirement of a macroscopic view of the music, but also inhibits performer interaction, an important component of the music, socially and aesthetically [19].

Other explicit considerations of polytempo music are sparse, and the paucity of published research in this area is marked by the writings of lamenting composers desperate to explore more advanced musical timings, such as the late Stockhausen [17]. A useful website, run by artist John Greschak [8], contains more information and unpublished articles about polytempo, as well as an annotated list of polytempo music. To our knowledge, there has been no previously published work in the area of music interaction or interface design that has significantly addressed musical tempo as the focus of control, nor explicitly considered the composer and composition as target user or task.

3. A SYSTEM FOR POLYTEMPI

There are two principal requirements of a system allowing composers to interact with tempo and polytempi: a representation of the polytempi, including the temporal relations between parts (the notation); and a method of manipulating and managing the tempo of such parts (the system). In this latter case, interaction should occur in realtime, in order to quickly allow the auditioning of alternative material and making of expressive refinements.

However, before further considering issues of system design and implementation, we must tackle one of the fundamental goals of our research: the design of a notation for polytempi, upon which the system will be based.

3.1 Notation

The design of our system was arrived at by drawing on our prior research into notations for performance and composition in music and other expressive arts [2][5]. The lack of previous work in this specific problem encouraged us to look for analogies in other disciplines and fields where it is necessary to handle parallel streams, signals and processes – such as physics, data communication, computer security, graphics, and engineering fields.

In facilitated cross-disciplinary meetings with 10 different specialist research groups (see Section 8, for a full list), the concept of phase and synchronization was highlighted in a number of non-musical activities, possibly the closest cousin of which is physical sound itself. A periodic waveform, such as a sine wave, at any moment has a phase, frequency and wavelength that might be adapted to music, in the forms of musical position, tempo and bar length, respectively[1].

Considering a musical part as a periodic signal, the challenge moves to representing multiple signals so that the relationship between them is evident. In many fields, phase can be plotted or graphed as a function of other properties of a given system, such as time (e.g. phase plot) or frequency (e.g. Bode plot). In this manner, it would be possible to plot musical position on a vertical axis against absolute time on the horizontal, but this would only be useful in plotting absolute synchronization and absolute time offsets – tempo would be implicitly presented as line gradient, and relative alignments would also be difficult to identify. Instead, we propose a plot of the phase of one signal against the phase of another, as in Figure 2(a). In music, this is the musical position within the bar of one part, against that in another part, as shown in Figure 2(b).

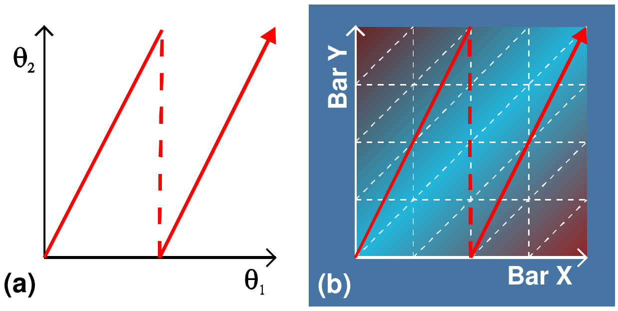

Figure 2. Plots of phase against phase.
(a) A general case. (b) An adaptation for musical purposes (4/4).

Although the plot no longer allows the reader to deduce the individual tempos of each part, the relationship between them is clear – a diagonal line (45 degrees) implies matched tempi; steeper or shallower and one part is faster or slower than the other. More importantly, the bar-level phase difference is also displayed, allowing the reader to easily deduce points of relative alignment, as shown by the guidelines in Figure 2(b). From the diagram it is possible to see the salient factors of the polytempi process – the relative phases and synchronization of two parts – and extrapolate how changes in each tempo, which affect the gradient of the plot, will affect the degree of synchrony over time. Figure 3 gives an illustration of a musical application.

To further illustrate how the plot functions, consider how Reich's "Piano Phase" would be represented: With two parts featuring close yet differing constant tempos, the line would be drawn with a gradient slightly off-diagonal. One part would reach the end of the bar sooner than the other, prompting the line to 'wrap-around' using the dashed lines, as in Figure 2. The wrap-around line illustrates the relation between the two parts'

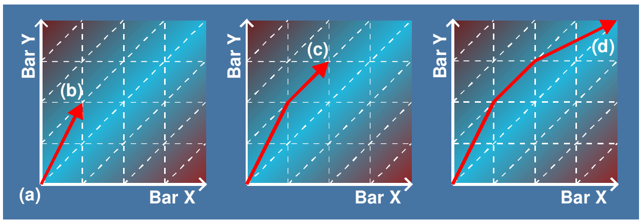

Figure 3. An idealised example of using a musical phase plot to manage polytempi.
(a) Part Y is progressing faster through the bar. (b) The part is slowed to the tempo of its counterpart, leaving them offset by 1 musical beat. (c) The part is again slowed so that, by (d), the parts are back in sync.

bar lines (the other part's formed by the axes of the graph), and gradually creeps across the grid, as the bar lines diverge, eventually converging on the opposite extreme of the bar, whereupon the process concludes, having regained synchrony, albeit a bar adrift.

3.2 Interface

The examples above demonstrate how the plot can be used to inspect temporal aspects of a piece but, in order to be of use to composers, a system must allow the viewer to affect the tempi – to draw the line themselves – and react to what they see and hear.

It would be possible to expose the relative synchronisation as a control parameter, but this would require the composer to first select a reference part to which the synchronisation would be relative, effectively restricting tempo variation to a single part, at any given time. Instead, we elected to simply control the tempi of both parts independently.

In addition to these two fundamental variables, we envisaged additional control parameters. Notably, the composer will, at different times, wish to affect tempo variations of varying scale. With Reich and Nancarrow, the tempo changes were gradual and finely-controlled, but other composers, such as Alejandro Viñao, desire the expressive freedom to make both fine and more abrupt, coarser variations. Thus, a third variable of control range (or resolution) is required. Finally, observing that temporal harmony involves the varying between two extremes (temporal consonance and dissonance), and that most pieces revolve around the journeys between them, we introduce a fourth factor in the interaction: a "gravitational" element that draws the two parts into consonant temporal congruity, to a varying degree. Altogether, this requires an interface offering at least 4 degrees of freedom, corresponding to: tempo of first part, tempo of second part, tempo control resolution and influence of gravity.

Our interface could simply be formed from common input widgets (sliders, rotary knobs, etc.). However, in designing our prototype, we turned to human gesture, where the body affords a large variety of motions to which our scales might be effectively mapped, and where their interrelationships and dependencies might be implicitly reflected. Gesture is often seen as a 'natural' interaction mode for computer-based musical applications, owing to the physical and tangible nature of interaction in traditional music making [13]. In this vein, we elected to use gestures, motions and actions that would not appear out of character with those established in live musical

[1] Amplitude, the remaining fundamental characteristic of audio signals constitutes an instantaneous property, and might be seen as the counterpart to similar musical properties such as dynamics, pitch, instrumentation, etc.

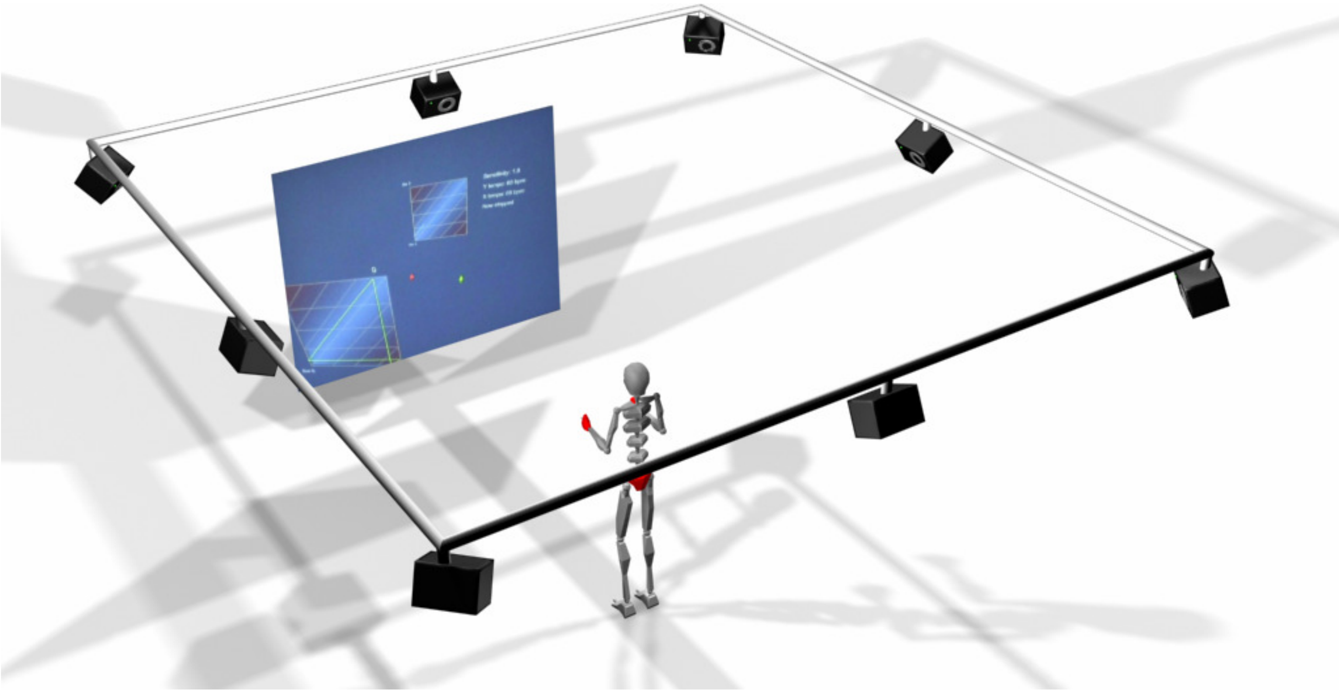

Figure 4. A Vicon™ Motion Capture-based system designed for controlling and representing polytempi.

performance. Specifically, the similarity of expressive roles between our user, the composer and that of a conductor was a significant influence in our selection. The intended result was a method of interacting that would make a user more comfortable in their manipulation of the system, where musical-like physical actions prompted clear musical results and users did not feel inhibited or self-conscious by having to make overt, overly-exuberant and uncharacteristic gestures.

Projecting the phase schematic onto a wall-mounted screen, a Vicon™ Motion Capture system [18] was used to capture body motion. A similar system has been used to control synthesizers and sound generation [6], but we could find no published account of an attempt to use such a system and gesture to allow higher-level, realtime control of musical composition and expression.

Our system (see Figures 4 and 5) was designed so that the height of each hand would set the tempo of each respective part. Walking forwards or backwards set the tempo range addressed by the hands, literally allowing more "up-close" adjustments or broader handling "from a distance". Appropriately, the effect of gravity could be controlled by bringing the hands closer together laterally, so that clasped hands (vertical and horizontal proximity) would ultimately bring about synchronisation of both tempo and relative position. Additional gestures were added to start and stop playback (a quick clap), and to allow the user to lock the tempo of each part (turning the respective palm up) so that they could focus on the other.

4. TECHNICAL DETAILS

The Vicon™ system works by using multiple cameras that can detect infra-red light reflected off small reflective balls attached to the subject. Belts, hats, gloves and suits adorned with these balls can be worn to allow untethered movement to be recorded within a confined space. The raw camera data is processed by the Vicon™ system into a realtime stream of 3D coordinates, which can be combined into groups representing different bodies and limbs.

For our system, we used two gloves and a belt to allow us to determine the position of the hands, relative to the body, and the position of the body relative to the space. The data was piped over a TCP/IP network to a PC running Cycling 74's Max/MSP and Jitter. Using C++, we developed a Max/MSP external that converted the data packets into usable Max variables. Variables corresponding to the positions, velocities and orientation of the waist and each hand were connected to the respective control variables of a MIDI playback engine (playing pre-recorded piano or percussive parts), so that they could be appropriately manipulated. In turn, the control variables, together with the status variables of the engine, were then passed to a Jitter patch that constructed a graphical representation of the musical phase plot, to be fed back to the screen.

Despite a diverse collection of protocols, the different technologies integrated well, and a basic system was up and running quickly, allowing us time to iteratively refine the interaction. The system outlined in Section 4 was designed to encapsulate the relative properties of the synchronisation between parts and, in doing so, would provide only limited insight into the more absolute characteristics of the performance – notably, absolute tempo or absolute part position. In Figure 4 and 5, the screen shows the musical phase plot in a 3D perspective, whereby a different plot is presented for each bar of a single part, flying forward in an abstract 3D space, appearing at a distance from the upper right (allowing bars to be read left-to-right), at a speed matching the part's tempo. The user is thus given the impression that they are progressing through the piece, and at what rate.

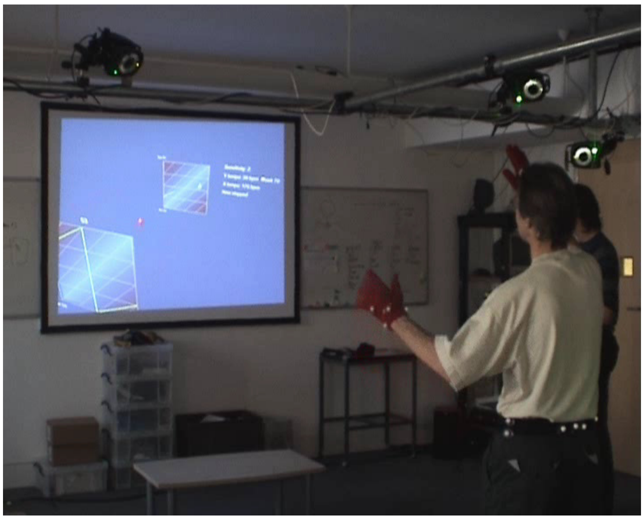

Figure 5. Alejandro Viñao (foreground) using the prototype.

5. DISCUSSION

Following development of the basic architecture, implementation of the prototype followed an iterative design process, based on feedback from our own interaction with the system, and three trials by Alejandro Viñao, which produced positive and useful feedback.

The difference between our interactions and that of a practising composer were revealing. To test the system, we used a variety of movements to ensure a robust and varied interaction. As demonstrated by the video (see Section 6), Alejandro's gestures were significantly more subtle, focusing on fine control – reiterating the utility of the resolution control.

For Alejandro, the strength of the design concept was already evident in the early version of the system we had available for his first visit. The forwards-backwards tempo control resolution feature was introduced for his second visit and refined in the third to allow a level of temporal control at which he could comfortably and confidently effect temporal manipulations in the music. Experimenting with temporality in a small selection of pre-prepared pieces, Alejandro mentioned that he was already eager to try the prototype with music of his own creation.

As with many musical instruments, mastering the interaction might have required more than the short exposure afforded Alejandro. In this respect, the "gravity" feature demonstrated its potential as a helper device, assisting actions that would otherwise require fine control and hand-eye coordination – allowing Alejandro to more easily target and achieve alignment and temporal consonance. The prototypes were only implemented with a basic gravity effect that brought the two parts closer together in absolute musical position. A more flexible feature, whereby the effect might gradually match tempos or align position relative to either the bar, beat or sub-beat, would further improve the usability and creative flexibility of the system.

Alejandro observed that the system was well-suited to inspecting, manipulating and adapting to polytempo processes operating at the level of the bar – either within a bar, or relative to the bar line. However, he noted that it was difficult to orientate oneself to the macroscopic aspects of the piece. For example, the system's difference in feedback between the two parts when offset by half-a-bar and when offset by one-and-a-half bars was minimal, yet could potentially hold important musical implications. The display of absolute positions and tempos was limited to the status readout on the right of the display, together with the appropriate labelling of the axes (indicating the current bar of each part). Similarly, although the extrusion of the plots into 3D succeeded in providing a sense of progress and time passing, the wrap-around lines now leapt from plot to plot, making it harder to observe bar-to-bar trends, and we were led to conclude that the original, static 2D design might be more suitable in most musical applications.

Furthermore, a better macroscopic impression would be afforded by adjusting a 2D musical phase plot (drawn relative to bar lengths), to one where the axes simply represent a continuing, absolute musical position within each part. The viewport would then pan over the current musical position appropriately. This would obviate the need for the wrap-around line, which might reduce the visibility of bar-to-bar relationships, but which could be replaced by appropriate annotations to make bar transitions and relationships more explicit.

6. CONCLUSIONS

This paper identified an area of musical expression that has received relatively little attention from technologists and music researchers. In an effort to tackle the barriers between composers and polytempi, we have proposed and tested both a notation for representing multiple tempi in music and a gestural system for interacting with them in realtime.

Alejandro Viñao's assessment demonstrated the aptness of our underlying design concept, while identifying a number of minor interaction issues that would be relatively easy to address. Our system, however, is but one possible solution to the problem, based on but one suggestion for a notation supporting polytempi. It is yet to be established how well our system scales to pieces with more than two differing tempi (i.e. using multiple plots or multiple axes). Furthermore, a major challenge will be the integration of polytempi notation with both live performance and other elements of music (melody, harmony, dynamics, rhythm, etc.), both of which would afford the composer or conductor greater possible creative freedom for realising music.

7. SUPPORTING MATERIAL

A computer animation demonstrating the system, including the supported gestures, as well as a video of Alejandro Viñao using the system is available online at:

http://polytempi.nashnet.co.uk

8. ACKNOWLEDGMENTS

Special thanks to Alejandro Viñao, who presented us the challenge and gave us invaluable insights into modern musical practice; and additionally to Tristram Bracey and Joe Osborne, who worked on developing the prototype. For a wide range of other insights, thanks to the various groups who kindly met with us: the Digital Technology, NETOS, OPERA, Rainbow and Theory & Semantics research groups here in the Computer Laboratory; the Inference group at the Cavendish Laboratory; the Signal Processing group in the Engineering Department;

and the Socio-Digital System team at the Microsoft Research Centre. For additional input thanks also to Ian Cross, of the Centre for Music & Science. Lastly, many thanks to the Leverhulme Trust and the Engineering, and Physical Science Research Council (EPSRC), without whose financial support this project would not have been possible.

9. REFERENCES

[1] Bellaviti, S. Perception, Reception, and All That Popular Music: An Interview with Alejandro Viñao. In *Discourses in Music, 6,* **2** (Spr-Sum'07), University of Toronto, 2007.

[2] Blackwell, A. and Collins, N. The programming language as a musical instrument. In *Proceedings of PPIG 2005* (Brighton, UK, June 29-July 1, 2005), 2005, 120-130.

[3] Clark, E. Rhythm and Timing in Music. In *The Psychology of Music (Second Edition, ed. Deutsch, D.)*, Elsevier Press, 1999, 725-792.

[4] Collins, N. Relating Superhuman Virtuosity to Human Performance. In *Proceedings of MAXIS* (Sheffield Hallam University, April 12-14, 2002), 2002.

[5] Delahunta, S., McGregor, W. and Blackwell, A. Transactables. *Performance Research Journal, 9,* **2** (Jun. 2004), 67-72.

[6] Dobrian, C. and Bevilaqua, F. Gestural Contol of Music: Using the Vicon 8 Motion Capture system. In *Proceedings of NIME'03* (Quebec, Canada, May 22-24), 2003, 161-3.

[7] Ghent, E. Programmed Signals to Performers: A New Compositional Resource. In *Perspectives of New Music, 6,* **1**, 1967, 96-106

[8] Greschak, J. *Polytempo Music Articles.* Available at http://www.greschak.com/polytempo. Last Updated: Jan. 15, 2008. Last Checked: Jan 30, 2008.

[9] Grout, D. and Palisca, C. *A History of Western Music (Fifth Edition).* W. W. Norton & Co. Inc., NY, 1996.

[10] Howard, D. and Angus, J. *Acoustics and Psychoacoustics (Third Edition).* Focal Press, Oxford, UK, 2006.

[11] Jordà, S. New Musical Interfaces and New Music-making Paradigms. In *Proceeding of New Interfaces for Musical Expression* (ACM CHI'01), ACM Press., New York, 2001.

[12] Ligeti, L. Beta Foly: Experiments with Tradition and Technology in West Africe. In *Leonardo Music Journal, 10,* 2000, 41-48.

[13] Magnusson, T. and Mendieta, E. The Acoustic, the Digital and the Body: A Survey on Musical Instruments. In *Proceedings of NIME '07* (New York, June 6-10), 2007.

[14] Piston, W. and DeVoto, M. *Harmony (Fifth Edition).* W. W. Norton & Co. Inc., NY, 1987.

[15] Potter, K. *Four Minimalists: La Monte Young, Terry Riley, Steve Reich, Philip Glass (Music in the Twentieth Century)*, Cambridge University Press, Cambridge UK, 2000.

[16] Read, G. *Music Notation: A Manual of Modern Practice (Second Edition).* Taplinger Publishing Company, New York, NY, 1979.

[17] Stockhausen, K. How Time Passes. In *die Reihe, 3* (Musical Craftsmanship), 10-40.

[18] Vicon Motion Systems. *The Vicon MX Motion Capture System.* Detailed at http://www.vicon.com. Last Updated: Jan. 30, 2008. Last Checked: Jan 30, 2008.

[19] Williamon, Aaron. *Musical Excellence: Strategies and techniques to enhance performance.* Oxford University Press, Oxford, UK, 2004.