

Weak Interactions, Strong Bonds: Live electronics as a complex system

ODED BEN-TAL

Kingston University, Kingston, UK
Email: O.Ben-Tal@kingston.ac.uk

This article examines works for live, interactive electronics from the perspective of complex dynamic systems, placing the human–computer interaction within a wider set of relationships. From this perspective, composing equates to constructing a complex system with the performer(s) and the computer as key players within a wider network of interdependence. Using the author’s own compositions as examples, this article investigates the utility of a system view on interactive, live electronics.

1. INTRODUCTION

When live, interactive electronic music was reincarnated in digital formats starting in the 1980s, a major motivation (or at least selling point) was to release performers from the tyranny of the tape.¹ Having to synchronise with a fixed soundtrack limited performers’ expressive abilities. The hope was that computers would be able to follow human actions, freeing performers to become more expressive and play in a more natural way. It could be argued that, in reality, the main difference was a shift of responsibility for keeping the human and electronic components of a piece together: from the performer on stage to the computer operator off stage. In the more successful pieces, performers are able to inhabit their modified musical roles with ease and incorporate the extra dimension of electronics into the performance, even though they can only attain a limited understanding of the operation of the electronics and therefore why things happen the way they do. However, in many pieces, for the live electronics to work properly according to the composer’s plan, so many adjustments are required in real time that performers just give up on understanding the relation between what they do and the final musical outcome. In other words, the effective integration of performer and electronics in mixed pieces is a *compositional* challenge much more than a technological one. Well composed pieces using *fixed* sounds – some of Jonathan Harvey’s pieces come to mind – integrate the acoustic and electronic components effectively while allowing performers the ability

to perform expressively within the constraints of the score and the fixed element. It is this relationship between the acoustic performer and the electronics that is a central question in live, interactive music.

Any composer deciding to use live electronics has to acknowledge the fragility of this field. Not only is there the always present possibility that the electronics will not function properly in performance, but also rapidly changing technology threatens to make any piece obsolete (Bonardi and Barthélemy 2008). So perhaps a prudent first question to ask before beginning any composition of this kind is *why?* Or, more precisely, *is the live element necessary?* At some level, the moment you route the microphone signal through an effect (e.g., delay) the answer to that fundamental question would be yes. But, at least for me, there needs to be a stronger imperative or a more specific need that *requires* the electronics to operate on-the-fly in real time for the realisation of the piece. It is this thinking that led me to focus on mutual listening between the performer and the machine as the central driving force in the pieces I compose: the computer extracts data from the performer’s sound, evaluates it and responds appropriately; similarly, the performer is expected to listen to the electronic sounds and respond to them.

The terms ‘live’ and ‘interactive’ in electronic music are not without contention (e.g., Stroppa 1999; Emmerson 2017) and even the very definitions of what counts as live or interactive are up for debate (Drummond 2009). My concern in this article is with pieces mixing acoustic instruments (including the voice) with electronic sounds, at least some of which are produced live in performance. Examining the relationship between the performers and the computer in such a scenario, Rowe (1996) proposes two prototypical approaches: the Instrumental Paradigm at one end and the Player Paradigm at the other end. According to Rowe, the Instrumental Paradigm is concerned mainly with ‘constructing an extended musical instrument’ (Rowe 1996); the electronics become an extension of the instrument, with the same fluency and immediacy of control and response (Croft 2007). There are echoes of Csikszentmihalyi’s concept of flow in there: merging of action and awareness, fluently

¹Also see Risset (1999) who offers a more nuanced view about the rationale for his foray into this field.

exerting control, matching skill and challenge (Csikszentmihalyi 1996).

The Player Paradigm, at the other end of a continuum, is focused on constructing an ‘artificial player’ (Rowe 1996); for example, by implementing compositional processes computationally. George Lewis’s *Voyager* fits within this paradigm: the computer generates symbolic sequences based on internally defined procedures as well as data from the performers; the sound production is outsourced – for example, to a Disklavier (Lewis 2000). Shimon the robot is another example aiming to create an artificial player that can perform, compose and improvise music, including in partnership with human performers (Hoffman and Weinberg 2010; Savery, Zahray and Weinberg 2021).

In between these two extremes – an autonomous performer on the one end and an extended instrument on the other – there are many possibilities for systems that are concerned to varying degrees with both instrumental coupling and compositional processes. Hsu (2010) describes an improvising system that is concerned with timbre and sound production and was developed in collaboration with one main performer; this approach is different from *Voyager* or Shimon in its specificity – aspects of the electronics are tailored to cooperate with specific instruments. Utilising the same concept with different instrumentation entails adapting the system. The examples of my own music, which will be described later, fall even closer to the instrumental end of Rowe’s spectrum, while still retaining concern with compositional processes and a degree of autonomy.

Rather than focus exclusively on the human–computer relationship in these pieces, I propose to analyse them as complex systems and place the human–computer interaction within a wider set of relationships at play in the performance as well as during the composition/design process that gave birth to this complex system.

1.1. Three perspectives on complex systems

Working on vision processing, Marr (1982) analyses complex systems into three levels. At the top is the computational level that addresses the overall function or goal of the system. This level is underpinned by an algorithmic level: the processes that enable this goal and the data processing that happens there. The lowest level is that of the implementation.

Commenting on Marr’s work, McClamrock (1991) observes that complex systems may have more than three levels of organisation and that Marr’s framework conflates structural and functional decomposition of the system. Instead, he proposes that the three levels are better viewed as three perspectives or types of questions we can ask about a complex system (or any of its constituent

components). This separates the structural analysis, which can locate any number of structures and hierarchical levels within a complex system, from a functional analysis that asks questions about those structures. The first of his three perspectives – format and algorithm – is focused on the purpose of the computation, or in McClamrock’s own words: ‘What is the program?’ The second perspective – content, function and interpretation – looks at how the computation is achieved through a structural and functional analysis, while the third concerns the implementation of the algorithm – for example, is it in hardware? Embedded into this framework is the idea that complex systems need to be examined in relation to contexts. Specifying what computational task is accomplished by a component is done in relation to the aims of the system. Identifying how this task is achieved requires identifying how components interact.

Figure 1 illustrates a performance of a live, interactive piece as a complex system. This is a heterogeneous network particularly from the perspective of computation: the electronics explicitly process information; the human performer does process information but accomplishes much else besides; the instrument takes inputs, transforms them and produces outputs, but we rarely consider instruments as information processing units; the score is a representation and thus encapsulates information – it is a communication channel that imposes constraints rather than actively transform information; and, a MIDI interface device performs a similar function though at a much more basic level. Looking at the way information is transmitted and transformed within the system will provide insights even if the analogy to a computational system is imperfect.

Both the performer and the instrument are themselves complex systems and the computing infrastructure that underpins the code has a few layers between human programming and the hardware. However, the following discussion will have to be limited to salient aspects of the interactions at, and close to, the top organisational level.

What types of information are passed between the different components? There is a strong coupling between performers and their instruments that includes sonic, tactile and visual aspects. In many cases the performer will also be operating an interface device that sends electronic (commonly MIDI) signals to the electronics. The connection between the performer and the interface is visual and tactile. The input to the electronics is primarily the audio signal from the instrument and the electronic signal from the interface. The performer hears the audio produced by the electronics and (hopefully) follows the score, but may have also modified some aspects of it, usually

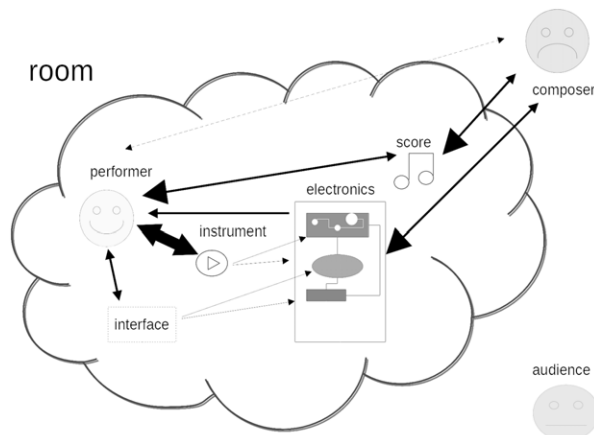


Figure 1. Performing a live electronics piece envisioned as a complex system.

in discussion with the composer during the rehearsal process.

The audience, the composer and the room are all important parts of the context or environment within which the complex system operates. One can argue for including the audience as a component of the system – communication between performers and audience is not one sided but interactive in some ways. However, with the rare exception where audiences are involved explicitly (e.g., where data is collected from the audience), the audience provides an external context to the system; performers ‘read’ the audience during the performance and react in subtle ways to the shared experience created between the performers and the audience. Thus, performers adapt their playing in the moment to the audience just as they adapt to different spaces. Both the audience and the space, therefore, can be considered part of the external context – an environment within which the system operates. The composer influences the performance indirectly, having shaped the score and the electronics; but, as we will see, this process is not entirely one sided. The choice of notation and the implementation of the electronics constrain and steer the composition process.

In the remainder of this section, I will describe key aspects of live, interactive electronics with reference to my own work. In Section 2 I will examine these as a complex system, while Section 3 will provide a summary and conclusions.

1.2. Music information retrieval

We use the metaphor *machine listening* to describe the process of converting audio signals into numeric descriptors that a machine can further process and evaluate. As Collins (2011) points out, despite all the advances in MIR techniques, there is still a

significant gap between this data conversion process and human listening. Furthermore, seemingly efficient MIR approaches may be un-musical (Sturm 2017). For machine listening to be part of a musical interaction with human performers it needs to relate to human musical concepts.

My first piece of live, interactive electronics – *Anemoi* (2004) for solo flute – uses pitch tracking and onset detection as factors in determining the electronic sound. Later pieces, particularly after I switched from using pure data (Pd) to SuperCollider, add additional techniques including extracting the strongest partials, timbral features such as spectral centroid (roughly the centre of mass of the spectrum), spectral entropy (‘peakiness’ – a rough estimate of periodicity/noise), or the width of the spectral distribution, as well as auditory features such as perceptual loudness and sensory dissonance. There is some use of direct mapping of these features onto synthesis parameters, but mostly these data are linked with decision making within the electronics (discussed in Section 1.4).

1.3. Modularity

The electronics for all the pieces discussed here consist of independent modules, each of them interacting directly with the performer. From the system point of view (see Figure 1) the top-level electronics is more of a hub or container rather than a controlling layer (except in one piece, see next section). In *Anemoi* there are four modules corresponding to the four movements of the piece. *Metaphors of Space and of Time* (2015) similarly consists of four movements but the electronics include six modules. Each movement is paired with a primary electronic process (except the third movement which has two); but the performer can mix the six modules live via a set of foot pedals. In other words, more than one module can be active at any given time and the player has freedom to choose which ones to activate (including the option of none).

Zaum: Beyond Mind (2010–13) is a sound theatre piece developed with fellow composer/performer Caroline Wilkins (Ben-Tal and Wilkins 2013). The electronics include multiple, independent components that I mix live on stage. These components fall into three broad categories: (1) direct audio processing, (2) fixed soundfiles, and (3) modules that use machine listening to selectively respond to some sounds from Caroline’s performance. In a performance,² one active module collects melodic contours, adding detected notes to a list until a large gap in pitch or time is discovered; this gap is used as a marker that the gesture is done, at which point the contour list is mapped onto

²Extract from *Zaum* at the Sonorities Festival in Belfast (starting 6 minutes into the video): https://youtu.be/_rqr58OP0jc?t=364 (accessed 25 November 2022).

subtractive synthesis, resulting in a transfer of material from the performer to the computer. A separate module detects loud and sharp sounds (technically a sharp decay of amplitude at the end of the sound) and triggers interjections of bandoneon samples – in other sections of the performance, Caroline plays the bandoneon; the samples heard here were recorded during the development process. A third module listens for unpitched sounds (whispers) and responds with filtered noise sounds (heard around 6:35). At the same time we can hear, in the background, one of the fixed soundfiles used in the performance: soft sustained (but subtly changing) sounds. Finally, there is some direct processing in the form of ring-modulation and reverb applied to Caroline's voice.

1.4. Decision making

Many of the modules in each piece link machine listening to decision making. This decision making is mostly formulated as binary choices. In the final movement of *Anemoui*, the computer uses onset detection and pitch tracking to sort the music into two categories: short rapid notes vs. sustained notes. When the average of the three recent inter-onset intervals (IOI) is below a threshold the computer flags this as short rapid notes. Sustained notes, on the other hand, are identified through small variation in pitch tracking and no onsets over a set duration (approximately 0.6 seconds). This binary choice is used to determine the response from the electronics, which in this instance respond in kind: long or short notes matching what the flute played. This suggests a simple imitative relationship, but in reality the unreliability inherent in the feature extraction leads to the occasional incorrect identification. As a result, the electronics sometimes add sustained notes over the flute's fast ones or vice versa.³ Computationally this is an error but musically it is not – it adds moments of contested relationship in an otherwise conformant context, to borrow Nicholas Cook's terminology for audiovisual relationships (Cook 2000).

Another form of binary choice is a true/false selection: does the input match certain characteristics – for instance, the example from *Zaum* earlier where a module only responds to instances of non-pitched vocal sounds. Individual modules implement a simplistic form of listening but when active in parallel they produce a rich and complex musical fabric. This, of course, is a hallmark of complex systems: emergent behaviour from the combination of simpler parts. At the same time, this selective response in the electronics (selecting whether to respond or not and selecting how to respond) leads to a less

straightforward correlation between the sound from the performer and the electronics.

Sometimes, machine listening is used to steer algorithmically defined processes. One of the modules in *Present Perfect* (2017) generates a melodic line in real time. The initial pitch is taken from the cello and subsequent pitches are chosen from a limited set of intervals from the previous note, but the current cello note exerts some pull; these two forces are combined using a weighted random choice. In other words the musical logic is based on both horizontal and vertical interval relationships. And, like other forms of contrapuntal writing, the lines are shaped by both internal factors and relationship to other lines. The shape of the cello melody and the algorithm shaping the contrapuntal line are interlinked; I arrived at their final form gradually as both were developed concurrently.⁴

A recent piece was my first in which decision making was implemented at the top level of the electronics. *One, Two, Many* (2021) is scored for two flutes (one and two) and electronics (the many), which still consist of independent modules. However, while in earlier pieces these modules were controlled manually (via MIDI controllers), here it is the computer that turns these on and off during the performance (or at least should). The computer assumes the role of an impatient listener seeking novelty. The first stage involves extracting features from the flute players: loudness, spectral centroid, spectral entropy and a chromagram.⁵ In the next stage, the computer tries to estimate stability, that is, change in the incoming signal. With each evaluation, lack of change increases a 'boredom' parameter, while change reduces it. When this 'boredom' meter crosses above a threshold, the computer changes things by either turning on an additional processing module or turning one off.

The stability and change are evaluated at three time scales. At the smallest scale, Kalman filters (Ribeiro 2004) evaluate the predictability of the spectral features: high predictability is taken as an indication of stability and therefore increases the 'boredom' meter. Two further evaluations of stability compare chromagrams summed over 10 and 30 second spans. Cosine distance is used as a similarity measure and the same 'boredom' meter is increased when successive spans are judged to be similar.

1.5. Mutual listening

So far I have described how the electronics react to the performer by selectively responding to particular elements and using information about the acoustic

³This can be heard in the recording available at: <https://soundcloud.com/odedbental/eurosjulian> (accessed 25 November 2022).

⁴<https://youtu.be/34zVcqXGHdk?t=533> (accessed 25 November 2022).

⁵Cumulative energy across the entire spectrum in each note. Chromagrams are often used to estimate chord or key from the music.

Figure 2. The end of the ‘Boreas’ movement from *Anemoi* with improvisation moments; open grey rectangles include instructions for pitch content and dynamic range with overall duration.

sound to control the electronics. However, I am interested in setting up a mutual listening scenario where the performer listens to the electronics and responds as well. The following examples illustrate several different ways this listening loop operates.

In the first movement of *Anemoi*, the flute sound is distorted (using ring modulation) and then sent into a set of feedback delay loops. Pitch extracted from the flute controls the amount of feedback: high notes increase feedback while low notes decrease it. The performer is tasked with playing a dangerous game: they have to locate the balance point where the feedback is enough to create a slow crescendo but avoid the inevitable explosion of too much feedback. The score leaves sections for improvisation where the player can choose high and low notes at will (see Figure 2).⁶

Non Sequitur (2015) also asks the performer to explicitly listen to the electronics. The piece uses a set of sensors,⁷ installed under the keyboard of a normal piano, to collect MIDI data while the player plays on the normal keyboard. These sensors are used to superimpose digital synthesisers onto the piano. One of the synthesisers generates a steady pulse, the pitch of which is a microtonal interval away from the pressed key, while the velocity is mapped to the pulse rate. The pianist has to listen to the sound and use the resulting pulse as the tempo for the next phrase.

In the first movement of *Metaphors of Space and of Time* (titled ‘Points’), the primary electronic process is based on delays, but the delay time is controlled by an estimation of the IOI of recent notes from the trombone. There are two delays (without feedback): the delay time of the first is set as the estimated average IOI, the delay time of the second is set as the reciprocal

of the IOI (1/IOI in seconds). The result is that one delay follows the player while the other counteracts – when the player slows down, the 1/IOI delay will speed up, almost like the computer is trying to trip up the performer. One feature of this combination – two delays at reciprocal time values – is that if the performer manages to play notes at exactly metronome mark of 60 (IOI of 1 second), all three sound streams – trombone and each delay – should synchronise perfectly. Torbjörn Hultmark ended up using this feature to structure his performance of this movement.⁸ This choice is not grounded in the notation but is an approach he developed gradually over the 15 or more times he performed this piece.

An extract from a performance of *Zaum: Beyond Mind*⁹ serves as another illustration. As noted earlier, at several moments in the performance I mix in fixed soundfiles. This happens ten seconds into the linked video when Caroline’s speaking voice comes through the speakers (starts with ‘Mind fragment’). Caroline reacts – visibly and theatrically, but later vocally as well. The result is a performed dialogue between Caroline on stage and ghostly duplicates of herself emanating from speakers around the room. The triggered soundfiles using her spoken words are my dramatised versions of the text, composed offline. The sonic sound material is Caroline, but layering, editing and subtle alterations of speed and pitch turn it into its own distinct contribution with which Caroline interacts. This points to the multiple interactions that operate in this performance beyond instrument/electronics. There is the sonic interaction between the electronics and Caroline’s acoustic sound, which is analysed and transformed by the computer;

⁶Recording available at <https://soundcloud.com/odedbental/boreasjulian>. (accessed 25 November 2022).

⁷PNOScan www.qrsmusic.com/PNOScan.php. (accessed 25 November 2022).

⁸As can be heard in this recording of the opening movement: https://youtu.be/RXG_euYtcY (accessed 25 November 2022).

⁹At the Logos Foundation – recording available at: www.youtube.com/watch?v=l-BicLttHjU. (accessed 25 November 2022).



The figure consists of four staves of musical notation for Cello, numbered 1, 4, 7, and 10. Each staff contains musical notation with various dynamics (mp, f, pp, mf) and performance instructions like 'gliss' and '3' (triplets).

- Staff 1: Starts with a half note G2, followed by a dynamic change from *mp* to *f* and then to *pp*. It features a triplet of eighth notes and another triplet of sixteenth notes.
- Staff 4: Starts with a half note G2, followed by a dynamic change from *f* to *pp*. It features a triplet of eighth notes and another triplet of sixteenth notes.
- Staff 7: Starts with a half note G2, followed by a dynamic change from *mf* to *pp*. It includes a glissando instruction and a triplet of eighth notes.
- Staff 10: Starts with a half note G2, followed by a dynamic change from *mp* to *f* and then to *pp*. It includes a glissando instruction and a triplet of eighth notes.

Figure 3. Sketch material for *Present Perfect*.

the modularity means that this level of interaction is applied separately to her vocal and instrumental performances (on the bandoneon). There is also a level of interaction between Caroline's onstage persona and the disembodied performance projected into the space from the speakers.

1.6. Composition process

A major constraint when working with machine listening, especially in real time, is the problem of reliability. Pitch tracking works fairly well for monophonic instruments but octave errors are still common. Extended techniques, chords and even vibrato in some cases will result in unpredictable results and any noise or interferences from other instruments will increase uncertainty. Even distinguishing between signal and silence or extraneous noise is not foolproof. As listeners, we adapt our concept of silence – which is never really silent – to the listening situation. Adjusting detection thresholds in feature extraction methods can somewhat compensate for different circumstances but not fully. Various spectral features are even more variable depending on the microphone, its placement, the instrument and the room. Composing these pieces, therefore, started with a period of experimentation that allowed me to discover what works and what I can achieve within the constraints of available tools

(e.g., MIR techniques) and my programming abilities. In other words, a research process: ideas and concept come into focus and initial plans are adapted as I discover unforeseen hurdles as well as opportunities.

The score and the electronics are intimately linked and developed in tandem and often benefited from working with the performers. Players were able to record some sketches and I used these recordings when developing the electronics. The sketches range from basic samples (e.g., the same note played with different mutes), to short phrases, to draft versions of sections from the piece. Figure 3 is taken from sketches I sent cellist Matthew Barley while composing *Present Perfect*, while Figure 4 shows the final version. The basic concept remained the same – alternating between a low E flat and activity in the high register – except that the material in the high register became isolated, fragmented figures rather than short melodies (and the low note arrives later). This change partly followed the decision that the piece should not have a clear start, an idea that came fairly late in the composition process, but it also reflects an element in the electronics: the computer latches onto sustained notes and continues them with a synthesised note using frequency-domain grating (Hartmann 1985). The computer is much less likely to identify the new, fragmented material as sustained notes compared with the melodic material in the sketch. Enhancing this contrast makes the system

Drift into the 'piece' without a defined start.

Prepare to perform: adjust chair, check bow, tap body of cello, play a few short, isolated notes, etc.

Start playback of soundfile, very quietly

Continue preparing: tune, move stand, check the computer is working, play a few notes (e.g. from the opening bars), etc.

increase volume of soundfile playback but still soft

Figure 4. The opening section of *Present Perfect* in the final score.

more robust and reliable. At the same time, the undulating nature of the frequency-domain grating – the overtones pulse in and out in this tone – bridges the gap between the low cello note and the higher material that now hovers around overtones of that same low E flat.¹⁰

We saw in several examples so far how the machine listening utilised in these pieces includes processing of the raw MIR features to infer gestural information, such as the melodic contours collected in *Zaum* (Section 1.3). This is another meeting point between the way I develop the electronics and my more general compositional thinking in which musical gestures (Ben-Tal 2012) are an important aspect. As noted throughout this section, in each piece the different components are specific – the aim is not to develop a general improvising machine. The exploratory composition process yields interlocking score and electronics that are tailor-made to the piece and the instrument(s) and often developed in collaboration with performers.

In some cases, I was also able to experiment with the performers in the studio and their comments informed the final stages of the composition. For example, in the final movement of *Metaphors of Space and of Time*, I

ask the trombonist to produce unpitched sounds and the electronics respond with similar short percussive sounds.¹¹ After a session trying out the piece, Torbjörn Hultmark remarked that he has to work hard to drive the electronics and it would be useful to have some longer electronic sounds occasionally to allow him momentary rests. The electronics were modified to generate the occasional longer noisy sounds, not just short percussive bursts.

1.7. Notation and improvisation

I view the notation as an approximation of 'The Piece' that, together with the other components in the system, elicits, guides and constrains the performer's creative freedom in realising it. One of the main challenges in enabling the mutual listening scenarios in these pieces is to communicate my own musical ideas while giving performers freedom. Figure 2 illustrates one method: short spans where the player is free to improvise; the score specifies a few parameters but leaves the parameter that controls the electronics free (recalling, in the Figure 2 example, that high/low notes control the amount of feedback). The score

¹⁰Recording available at: <https://youtu.be/34zVcqXGHdk?t=70> (accessed 25 November 2022).

¹¹Final movement: https://youtu.be/RXG__euYtcY?t=383 (accessed 25 November 2022).

(a) **Surface**

Note type and slurs indicates heirarchy of importance and pitch relationship. Can play pitch in any octave.
Each 'measure' is a harmonic field, improvise your way through these fields in the given order.

(b)

Figure 5. Two extracts from *Metaphors of Space and of Time*: a) is the opening of the second movement, specifying harmonic material only: b) is an extract from the fourth movement specifying mostly rhythmic information.

also instructs the player to base the improvisation on preceding material, thus giving them a framework or prompt to develop their ideas. The score for *Metaphors of Space and of Time* gives the performer even more freedom. In the second movement ('Surface') the score only specifies pitch: a harmonic skeleton (Figure 5a) of two alternating chords with some variation in each iteration (compare bars 1 and 3 with bars 2 and 4). The notation includes a hierarchy of importance (using duration values – longer is more important) and connections (slurs). The manner in which to articulate this is left open. The final movement ('Volume') asks the player to produce unpitched noise sounds. The rhythm and the tempo (very fast) are specified, and the score also asks the player to utilise different sound qualities (Figure 5b).¹²

These examples show how the notation guides the player, but the interpretation is also guided by the matching sounds coming out of the electronics. Each of the movements of *Metaphors of Space and of Time* has a distinct character that is the result of careful calibration of the tripartite relationship between the electronics, the performer and the score. Both the cyclic nature of the chords and the way the electronics sustain the trombone sound in this movement contribute equally to Torbjörn's choice to

perform 'Surface' as a slow movement. Similarly, the erratic noise bursts that the electronics contribute in the final movement mirror and reinforce the 'Manic' character specified at the top.

The final illustration comes from *One, Two, Many*, where the performers make simple choices within an otherwise notated context (Figure 6). The score includes some short optional figures that the players are free to play or to skip and the same applies to the repeat signs: player choose whether to repeat or not and can also repeat more than once. In a few places the score includes two alternatives and the player needs to choose which one to perform.

2. AS A COMPLEX SYSTEM

As we saw in Section 1, McClamrock (1991) proposes to interrogate information processing systems from three perspectives. The first focuses on the purpose – what is the goal? What is the information that is being processed and how is it represented? The remaining perspectives are focused on how is this goal achieved: how is the computational task organised and how is it implemented? The set of three questions can be applied at different levels of structural decomposition. While acknowledging that music is not information (or not just information), these three perspectives can nevertheless be relevant as a means of

¹²Recording available at: https://youtu.be/RXG__euYtcY (accessed 25 November 2022).

Flute1

Figure 6. Extract from *One, Two, Many*. Players can choose whether to repeat material within the repeat signs (top line) and how many repetitions. In the second line they choose whether to play the top or bottom staff. Players can also choose whether to play or skip the figures marked optional (last two lines).

interrogating the complex phenomena that music is – particularly in the case of live, interactive electronics, where the computing side *is* purely information processing. The discussion that follows will focus on the first two sets of questions and not on the more technical question of implementation.

Starting with the system as a whole (Figure 1), its purpose is to perform ‘the piece itself’ – transforming an abstract concept into concrete sound.¹³ This abstract concept is, of course, much debated. I do not wish to make a dogmatic argument for the existence of a singular one true ‘piece’, rather I make a pragmatic observation that composers, performers and listeners do have a concept of a piece that is related but not identical to any individual performance of it.¹⁴ Viewed as a complex system, the act of composing

a piece equates to constructing this system: directly shaping some elements – the score, the electronics – while others are just selected for inclusion – acoustic instruments and (hopefully) good performers.

Within the complex system, only performers have conscious access to the computational purpose of the system: enacting ‘the piece’. The implication is that only the performer can be said to have a representation relevant to the algorithmic task (recalling that part of McClamrock’s first perspective asks how the information is represented). Here we do not mean a purely mental representation confined inside the performer’s mind; rather, this is a perspective on a *system* that includes actions, objects and relationships, therefore the performer’s representation of the piece is embodied, embedded, enacted and extended (van der Schyff, Schiavio, Walton, Velardo and Chemero 2018).

Shifting down one level, we can view the score as analogous to a ‘programme’ that the performer runs within the system. Because human performers,

¹³I deliberately use the word ‘sound’ and not ‘music’ here. Music is what listeners make out of sound (Reybrouck 2020), and as stated earlier, I place those listeners outside the system itself.

¹⁴See also Marsden (2016) for further discussion.

unlike computers, are actually intelligent and also experienced (in music, culture and the world), this programme is not a closed set of instructions in the manner of a computer programme; performers do not merely execute the score, they interpret it. They need to develop this interpretation in relation to the other elements in the system, particularly the unfamiliar element of the electronics. After several performances, when he was already familiar with *Metaphors of Space and of Time*, Torbjörn Hultmark observed that it is not a piece that can be performed after two rehearsals; it takes time to discover the possibilities inherent within this complex system and to determine how he, as the most knowledgeable and capable element within the system, can steer it.

As described in Section 1.4, the electronics in all these pieces try to match what the performer does against defined categories. This, effectively, creates a rudimentary form of musical representation in the electronics. This representation is tied very specifically to each piece – the electronics developed for these pieces are not universal improvisation tools and are not modelling human forms of listening. Linking the electronics to higher structural levels – gestures or patterns that span longer durations – creates a distance between the performer and the computer (shifting away from the instrument paradigm towards the player paradigm). However, this distance opens a space for musical dialogue based on mutual listening and on an affinity between the human performer and the electronics. This affinity is manifested through a combination of sound parameters (timbral aspects below note level) and higher level characteristics of the music above the note/event level.

The computational task of the electronics is mostly delegated down one level to the independent modules. The functionality achieved at the top level is mostly restricted to evaluating whether the current input from the instrument is sound or silence and some reverb and compression applied at the end of the signal chain. Furthermore, in most cases a human hand turns the modules within the system on and off; the electronics are not an independent musical agent in that sense. The implementation of an overall listening strategy in *One, Two, Many* defines the computational goal of the system as a whole (something that was not well defined for the earlier pieces), resulting in electronics that are less dependent on human operation.

The individual modules within the electronics, on the other hand, are independent and do not require intervention beyond on/off and adjusting the output levels, that is, balance. They mostly implement dynamic processes that integrate internal logic with information extracted from the performer. Many incorporate recent history (of their own internal states and/or of the input from the performer) into the

computation. This means that the output from these modules is not directly correlated with the input from the player – very similar action from the player can result in different outcomes depending on the context. The result is live electronics that are predictable on a statistical level, but not on the level of individual sounds. Performers cannot predict what sound will come next nor when, but they do get to know the range of possible sounds and the kinds of textures and gestures they will hear.

3. SUMMARY

This article offers a view on live, interactive electronic pieces as complex systems, in which the interaction between performer and computer is part of several interlocking sets of relationships. Central to the analysis is a focus on the way information – broadly conceived – is transmitted and transformed within the system. The discussion is grounded in my own approach to live electronics, which hinges on mutual listening scenarios where the exchange of information between performer and electronics is explicitly designed into the piece. Composing the pieces means constructing the complex system through a research process where the different components are gradually developed in tandem.

The electronics themselves include both data processing and evaluation in the form of simple binary choices. Selectively responding to some sounds, or responding in different ways depending on the musical content of the signal, can give the appearance of musical intention. Torbjörn Hultmark once described performing with the electronics (with which he regularly improvises beyond performing *Metaphors of Space and of Time*) as being like ‘playing with a somewhat wilful partner’. This capacity of the electronics to produce modest surprise contributes to the sense that the electronics are semi-independent. Yet, the multifaceted affinity between the electronics and the performer – encompassing different aspects from timbre and pitch to gestures, patterns and performative elements – makes for a strong bond at the heart of a network of interactions.

From the perspective of my own composition practice, I see two main aspects for future development. Regarding the internal construction of the electronics, it will be interesting to expand beyond parallel, independent modules and to explore the musical possibilities in interacting elements. The binary choices in one module can influence other modules. Modules could also listen to each other and not just to the performer. A more ambitious project is to incorporate some planning into the electronics. One way of approaching this is to try to integrate some representation of the state, and the dynamics, of the whole system.

While doing something like this for the very general case – a universal improviser – is a daunting task indeed, tailoring it to a specific piece, while challenging, is more realisable.

REFERENCES

- Ben-Tal, O. 2012. Characterising Musical Gestures. *Musicae Scientiae* **16**(3): 247–61.
- Ben-Tal, O. and Wilkins, C. 2013. Improvisation as a Creative Dialogue. *Perspectives of New Music* **51**(1): 21–39.
- Bonardi, A. and Barthélemy, J. 2008. The Preservation, Emulation, Migration, and Virtualization of Live Electronics for Performing Arts: An Overview of Musical and Technical Issues. *Journal on Computing and Cultural Heritage (JOCCH)* **1**(1): 1–16.
- Collins, N. 2011. Machine Listening in SuperCollider. In S. Wilson, D. Cottle and N. Collins (eds.) *The SuperCollider Book*. Cambridge, MA: MIT Press, 439–62.
- Cook, N. 2000. *Analysing Musical Multimedia*. Oxford: Oxford University Press.
- Croft, J. 2007. Theses on Liveness. *Organised Sound* **12**(1): 59–66.
- Csikszentmihalyi, M. 1996. *Flow and the Psychology of Discovery and Invention*. New York: Harper Perennial Modern Classics.
- Drummond, J. 2009. Understanding Interactive Systems. *Organised Sound* **14**(2): 124–33.
- Emmerson, S. 2017. *Living Electronic Music*. London: Routledge.
- Hartmann, W. M. 1985. The frequency-domain grating. *The Journal of the Acoustical Society of America* **78**(4): 1421–5.
- Hoffman, G. and Weinberg, G. 2010. Shimon: An Interactive Improvisational Robotic Marimba Player. *CHI'10 Extended Abstracts on Human Factors in Computing Systems*, 3097–102.
- Hsu, W. 2010. Strategies for Managing Timbre and Interaction in Automatic Improvisation Systems. *Leonardo Music Journal* **20**: 33–9. https://doi.org/10.1162/LMJ_a_00010 (accessed 25 November 2022).
- Lewis, G. E. 2000. Too Many Notes: Computers, Complexity and Culture in Voyager. *Leonardo Music Journal* **10**: 33–9.
- Marr, D. 1982. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. Cambridge, MA: MIT Press.
- Marsden, A. 2016. Music Analysis by Computer: Ontology and Epistemology. In D. Meredith (ed.) *Computational Music Analysis*. Heidelberg: Springer, 3–28.
- McClamrock, R. 1991. Marr's Three Levels: A Re-Evaluation. *Minds and Machines* **1**(2): 185–96.
- Reybrouck, M. 2020. Music as Epistemic Construct: From Sonic Experience to Musical Sense-Making. *Leonardo Music Journal* **30**: 1–9.
- Ribeiro, M. I. 2004. Kalman and Extended Kalman Filters: Concept, Derivation and Properties. *Institute for Systems and Robotics* **43**: 46.
- Risset, J. C. 1999. Composing in real-time? *Contemporary Music Review* **18**(3): 31–9.
- Rowe, R. 1996. Incrementally Improving Interactive Music Systems. *Contemporary Music Review* **13**(2): 47–62.
- Savery, R., Zahray, L. and Weinberg, G. 2021. Shimon Sings-Robotic Musicianship Finds Its Voice. In E. R. Miranda (ed.) *Handbook of Artificial Intelligence for Music: Foundations, Advanced Approaches, and Developments for Creativity*. Springer International Publishing, 823–47.
- Stroppa, M. 1999. Live Electronics or... Live Music? Towards a Critique of Interaction. *Contemporary Music Review* **18**(3): 41–77.
- Sturm, B. L. 2017. The “Horse” Inside: Seeking Causes Behind the Behaviors of Music Content Analysis Systems. *Computers in Entertainment (CIE)* **14**(2): 1–32.
- van der Schyff, D., Schiavio, A., Walton, A., Velardo, V. and Chemero, A. 2018. Musical Creativity and the Embodied Mind: Exploring the Possibilities of 4E Cognition and Dynamical Systems Theory. *Music and Science*, 1. <https://doi.org/10.1177/2059204318792319> (accessed 25 November 2022).