

The Fragmented Orchestra

1. Introduction

The Fragmented Orchestra is a huge, distributed musical instrument, modelled on the firing of the human brain's neurons. This collaborative work, spanning music, art and science, has evolved from a fascination by its creators of the inherent rhythms and adaptive learning of spiking neurons. The piece is made up of one central space (installed, at the time of writing, at Liverpool's FACT gallery) and 24 geographically disparate sites across the UK.

Each of these sites has a "soundbox" unit installed containing a microphone, a speaker and a computer. The box streams live audio via the internet to the central installation whenever a threshold is reached; whenever a neuron fires. A firing event will cause fragments of human made and elemental sound to be relayed to the gallery and will also be communicated to the cortex as a whole. The combined sound of the 24 speakers at the gallery will be continuously transmitted back to the sites and to each of the 24 sites.

Together, these sites form a tiny 'cortex' of 24 artificial neurons, whose mechanism is adapted from a theory of spiking neurons developed by the computational neuroscientist Eugene Izhikevich.

The discrete audio streams are treated as individual inputs to the artificial spiking neural network; the series of impulses corresponding to each input's spiking activity is then used to granulate the stream's audio in real-time. The processed audio is then relayed to one of the speakers in the gallery. Finally, the aggregated mix of all 24 streams is relayed back to each of the sites, in order for visitors around the country to be able to hear and interact with the piece.

The geographical sites were selected for their sonic qualities, with the intention of building up a portrait of the UK's varied soundscapes in terms of geography and time scales. These range from inner city traffic, chanting from sports stadia, the hubbub of cattle auctions, and the chatter of migrating birds, combined with incidental and performed sounds from members of the public at concert venues such as Gloucester Cathedral and Belfast's Sonic Arts Research Centre. The public, invited to play the instrument at the 24 sites, will be able to hear the effect their playing has on the overall composition of the piece at each site and at FACT. As members of the public use the instrument, they become both player and audience of a vast and evolving musical composition extended across the UK.

For a global audience, the piece can be accessed via an interactive interface on the project's website [3]. A visitor to the site equipped with the Flash browser media plug-in



Figure 1. Photograph of the installation

can listen to a live neurally-processed stream from each of the individual sites, or combine audio from several sites to create their own mix using the material.

The Fragmented Orchestra was conceived by artist Jane Grant, physicist, musician and composer John Matthias and composer Nick Ryan. Network and server engineering were co-ordinated by Daniel Jones and Tim Hodgson respectively, and the project's visual identity was enabled by the London-based Kin Design. The Fragmented Orchestra was awarded the UK 2008 PRS Foundation New Music Award, whose funding part-enabled its realization.

This paper outlines and assesses a number of key aspects of The Fragmented Orchestra. Section 2 summarises the process of collaboration which underlined the entire project's life cycle; Section 3 describes the neuronal network at the core of the piece; Section 4 covers the architecture of the server that hosts the neuronal network and connects it to each of the broadcast sites; Section 5 outlines the network and communications infrastructure that hosts these sites, and the possibilities for interaction at the sites themselves; Section 6 describes the web interface; and Section 7 briefly examines the ancestry of the project. Finally, Section 8 assesses the outcome of the piece, and suggests directions for future work.

2. Interdisciplinary Collaboration

Given the complexity and breadth of the project's concepts, it was key to develop a working infrastructure of collaboration between individuals with different backgrounds. Throughout the year-long development period, artists, scientists, de-

signers and engineers worked in continual close proximity, exchanging and developing conceptual ideas. It was thus necessary at an early stage to establish a common lexicon of concepts that could be understood by all, with the design team familiarising themselves with the basic neuronal concepts, and engineers being briefed on the specific aesthetic requirements of the interaction design and sonic behaviours. This knowledge-sharing has historically been shown to be critical to a successful collaborative process [2].

After the initial, informal briefings, the tasks were divided into rough working groups: art direction, visual identity, software development, and hardware engineering. Some questions fell between these groups, and were tackled collaboratively as a discursive process. It was here important to create well-defined boundaries between working groups, with clear interfaces for programmers and hardware developers to work between; network protocols for client-server development had to be sufficiently specified for interoperability between networked components, and likewise for the various Audio Units of the server system. Similar boundaries occurred between the visual design team, the engineers creating the physical installation, and the developers working on the website, each drawing on work that had been delineated and performed by the others.

One fascinating consequence of working between complementary disciplines is the tangible effect that this has on the methods that each specialist adopts to do their work – for example, the composers behind *The Fragmented Orchestra*, unversed in engineering practices, quickly adopted the typical engineer’s approach of iterative development to realize the piece. This mode of transformative collaboration is identified by Prophet and d’Inverno [1] in their research into stem cell dynamics as “*transdisciplinary*”: “where the impact of disciplines on each other is strong enough to fundamentally affect the disciplines themselves” [1].

3. Neuronal Model

The majority of the processing of sensory information in our brains is thought to take place in the cerebral cortex, which contains a population of billions of neurons, each making thousands of synaptic connections with its neighbours. A single neuron can be thought of as a cell which generates a travelling spiking signal to these connected neighbours when the voltage on its membrane exceeds a certain threshold voltage, a process which is called ‘firing’. A neuron receiving several spikes simultaneously (or within a very small time-window) is likely to have its voltage pushed beyond the threshold level and will therefore in turn, send spike signals to its connected neighbours. Furthermore, connections between most neurons which cause spiking signals tend to become potentiated and those which do not become depleted, a phenomenon known as ‘Synaptic Plasticity’. The dynamics of millions of such adaptive, interconnected neurons thus provides an extremely rich behaviour,




Figure 2. Spike timing graph. x = time (ms), y = neuron index.

especially on a collective level of description (for a comprehensive introduction, see for example Gerstner and Kistler [5]) and patterns of firing regularly occur in groups of neurons. This is illustrated in Figure 2, taken from a mathematical simulation of a group of 1000 coupled neurons. For an example of such collective firing behaviour, see [4].

The neurons are numbered on the ordinate y-axis (with neuron number 1 at the bottom, and neuron number 1000 at the top) and the time, which runs from zero to 1000 milliseconds (or one second), is on the x-axis. This is therefore a simulation of one second’s activity of this group of 1000 artificial neurons. Every time a neuron fires a blue dot is placed on the graph at the appropriate time on a line horizontally drawn from that particular neuron. The dots on the graph can thus be regarded as firing ‘events’. In the particular graph shown, because of the plasticity of the neural connections, many of the events are centred in four bands, which appear as a pulse or ‘wave’ of spiking events in real-time.

The spiking events are indeterminate (not predictable in advance) but are certainly not randomly distributed, and, as is the case with the above scenario, can be very correlated. A rhythmic pattern such as the one pictured above is very likely to be connected with the ‘polychronous’ firing of a particular group of neurons [4], in which the firing of a particular neuron generates a sequence of events which stimulate a large number of neurons, which form a closed group with the connections between these neurons being reinforced through repeated firing of the first neuron (the group of neurons fire not in synchrony but with polychrony). It is the musicality of these sequences of firing events, or rhythms, which we are very interested in. In our research (see, for example, Miranda and Matthias [6]), Grant et. al. (2007), Matthias and Ryan [7]) we have started to look at what happens when each firing event is represented by a sonic event.

3.1. The Fragmented Orchestra

In The Fragmented Orchestra, each event is transformed into a short sample of sound streamed from one of the 24 sites. The cortex is tiny, consisting of just 24 neurons, but is sufficiently large to achieve many complex rhythms and collective firing behaviours. Each site is represented by its own neuron, which is also stimulated by the sound from that site. The samples of sound are relayed in real-time and last between 30 milliseconds and 2 seconds. We have adapted a mathematical model of biological neuronal networks developed by Eugene Izhikevich and others [4], which is one of many models of spiking neuronal networks which are based on a simplification of a model developed by Hodgkin and Huxley in the 1950s [8]. These models are essentially electrical in nature and consider the relationships between the flow of ions across cell membranes and the flow of voltage signals between the cells. In a sense, we have created a hybrid organism, involving 24 cortical neurons which are also all crude sensory neurons (they are stimulated by the 'volume' of the sound). This tiny organism is then distributed across the 24 sites and transformed into a musical instrument.

4. Server Architecture

The neural network itself is hosted by a Mac OS X application which resides on a Mac Pro server located at the installation site. The application takes the form of a graph of Audio Units, which handle the network streaming as well as the sound processing. The core of the audio processing is handled by two units, one containing the spiking neural network, and the other containing the granular sampler.

The network is based upon the Izhikevich model outlined above, taking its input from the physical network of audio sources plus its internal "noise", analogous to cortical neural activity. For each of these neurons, the unit sends two streams of data onwards through the chain, the first carrying the audio signal, the second containing spiking events generated by the neural network in response to both the direct stimulus to that neuron, and to activity elsewhere in the network.

The granular sampler is in many ways typical of this kind of unit: it selects 'grains' of sound from the input, of lengths varying from 30ms up to a second or more. These grains are output in a rearranged and layered form; the extent of this rearrangement being dependent on the chosen system settings.

The important difference with this granular sampler is that the timing and density of the grains is determined by the activity of the neural network, which in turn is partially determined by what's happening at the soundboxes ('partially' because the complex internal dynamics of the neural network mean that its behaviour is very far from a simple mapping of input to output).




Figure 3. Server architecture

Although the UI was not designed primarily as a "playing" interface, but rather a means of setting the initial conditions of a generative system, the server application's user interface nevertheless gives a high degree of control over the initial parameters of the system, and also allows ongoing adjustments while the system is running. It is possible, for instance, to alter the parameters of the spiking neuron model itself (its "recovery timescale" and "after-spike reset"). Similarly, the input to each neuron, corresponding to a single audio stream from a soundbox, can be scaled up or down independently of its actual volume, in real time, so that the channel's spiking activity is to some extent controllable, allowing each site to make a contribution to the overall output even if its average sound level is low. Each of the 24 input channels also has an adjustable delay buffer, from zero to 10 seconds, from which the sampler can draw grains. Thus fine adjustment of the degree of temporal "shuffling" of the source audio is easily achieved.

5. Soundbox network

The Fragmented Orchestra is fundamentally a distributed system, comprised of an interconnected network of communication nodes. Given that these sites are scattered throughout the length and breadth of the UK, it is entirely reliant on the availability of a network infrastructure that is capable of




Figure 4. Network architecture

transmitting audio data in real-time over a great distance. Indeed, this kind of project has only been rendered technically feasible in recent years, courtesy of the rapidly accelerating rate of consumer-grade internet connectivity.

The architecture of The Fragmented Orchestra's network is illustrated in Figure 4. The central server, which performs the neuro-granular processing, is located in the gallery space at FACT. This is connected via the internet to the 24 soundbox sites, each of which has a sufficiently fast broadband connection to relay audio streams in both directions in real-time.

Each of the soundbox sites hosts a lightweight computer, which has a multitude of parallel roles: to record audio data through its microphone, compress and transmit it over its internet connection, receive the processed audio from FACT, and output this sound through its Feonic Drive speaker. To make the program design as clear, efficient and modular as possible, each of these tasks is performed by an autonomous segment of the soundbox software, specializing in a single process and passing its output to the next segment; this is known as "multithreading" (Figure 5).

The project is thus distributive on multiple scales: it is comprised of a network of co-operative systems which are, in themselves, comprised of a network of components, all running in parallel to achieve their tasks.

This installation must be operational 24 hours a day, throughout its 3-month duration, so it is also vital that this software is as stable and resilient as possible. It must be able to cope with significant network delays and outages and must be configurable on a per-site basis according to specific needs; for example, one site must be scheduled to reduce its volume at night to avoid disturbing nearby residents. Moreover, due to the geographical distance, we must be able to comprehensively configure and control it remotely. Given these requirements, a piece of software to perform these tasks was written from scratch, developed over several months and




Figure 5. Soundbox process threads

thoroughly tested for reliability and stability.

This had a number of key features, most of which are configurable for each soundbox.

Network connection monitor. The state of the network connection is continually monitored by a dedicated thread, whose job it is to destroy stale connections and re-open the connection as rapidly as possible.

Data buffering. When the client first successfully makes its audio connection to the server, it prebuffers around 10 seconds of audio data. This ensures that a backlog of audio is always available in case of network jitter or lag.

Audio cache. The data buffer also serves a secondary role: if the network connection is lost altogether, the contents of this buffer are output from the speaker on a continual loop (and similarly at the network reception point at the server-side). Thus, some representative sound can continue to be played whilst awaiting the return of network connectivity.

Outage monitor. A system-wide scheduler is operational to reboot the computer in the worst-case scenario of an unrecoverable network outage or other nonfatal system failure.

The soundbox client is founded on a number of excellent open-source libraries (lame, mad, ogg vorbis and libshout), without which this part of the project would simply have not been possible. As a consequence, we will shortly be releasing the source code to the soundbox client back into

the public domain under a free license, to aid and encourage the development of future such projects.

5.1. Interactivity at the Sites

A sonic event at one of the sites will be fed into the neural server system; after a delay equal to twice the network buffer size, the output of the neural network will then be output from the loudspeaker at the site. Based on this feedback loop, a visitor to one of the sites can ‘play’ The Fragmented Orchestra as an interactive instrument, perceiving their actions as mediated through this chain of pathways.

This simple call-response behaviour itself seems to provoke a significant amount of engagement with a public audience; there is a great fascination in the process of creating a sound and hearing its echoes return, having travelled through a web of pathways spanning the nation. However, there is a greater degree of reactivity throughout the network than is immediately evident. Because of the interconnected nature of the neural network that governs the audio spike timings, audio events at one site can implicitly become associated with those at others, based on dynamic correlation patterns between the sites. As the system continues to run, and more of these connections evolve, this results in sound events at one site triggering those at another in synchrony – an audible example of the “pulse” or wave phenomena described in Section 3.

It is consequently possible to create an auditory cascade of events, the effect of which is akin to hearing echoes of sound from the recent past in distant corners of the country. In a learning system such as this, each event does not only have an immediate effect but has consequences for the long-term behaviour of the system.

6. Web Interface

For those unable to visit the sites that make up The Fragmented Orchestra’s physical architecture, the piece can be accessed using a web interface designed and developed for the purpose. The interface is an Adobe Flash application, based upon the sparse, geometric aesthetic that characterised the installation’s visual identity, and integrated with the rest of the website content. Indeed, the interface is the first content that the user is presented with, foregrounding the content of the piece above the supporting information that the remainder of the site provides; it was the intention to create an immediate, immersive experience, free of visual distractions and focusing the visitor on the sound of the work.

The experience itself is markedly different to that of visiting the exhibition site, though modelled on a similar concept. The user is presented with a network of nodes, each corresponding to a soundbox site, joined together by lines in a similar way to a diagrammatic neural network. A minimal set of control buttons enables the user to pause the piece, add and remove nodes, and switch to a full-screen display mode. No other visual elements are present, with the excep-

tion of a visual aura surrounding each node to indicate its sound state.

Consider the space that the nodes inhabit as a 2D plane. Each node conceptually outputs a live audio stream from its site to a limited area around it in this plane; the user then plays the role of a ‘listener’ within this space, whose position is determined by the mouse cursor. The amplitude and stereo panning of each of the nearby nodes is then modulated in real-time to create the impression of virtual sound-source positioning, with the analogy that the user is moving through the 2D space between this network of speakers – in much the same way that the visitor to the physical installation configures their experience of the work by moving between the speakers. This technique was tested and developed for AtomSwarm [11], an earlier performance piece by Daniel Jones, and was found to be a deceptively effective way to immerse the listener within the space of the composition itself.

Here, however, it is possible to add, remove and reposition nodes at will, so the user can select their own palette of sound sources from the array of varied sites. A choral performance at Gloucester Cathedral can thus be combined live with street sounds from Bristol, transforming The Fragmented Orchestra into an interactive instrument that can be used to create novel and unique compositions by members of the public.

It is critical to state that we view the website as complementary to the physical installation rather than as a strictly alternative method of experiencing the piece. Though efforts were made to optimize the immersiveness of the interface, it is a radically different experience to the cocoon-like physical immersion of the 24-channel surround system, whose embodiment excludes all other sensory stimuli to allow for complete focus on the sonic affects.

A common observation made by visitors to the installation is that the piece benefits from listening at various times of day at night. Given that the gallery is only ordinarily open during daylight hours, it is here that we feel that the 24/7 web access point succeeds in genuinely broadening the scope of the piece.

7. Origins of the Project

In 2004, John Matthias and Eduardo Miranda at the University of Plymouth developed the Neurogranular Sampler, an instrument which triggers grains of sound from sound files in a computer when artificial spiking neurons fire in an Izhikevich network [6]. These initial experiments formed a large part of the piece *Cortical Songs* written by John Matthias and Nick Ryan [7] which was released on Non-classical records in 2008. In this work, artificial spiking neurons control a set of lights, which flash when the neurons fire. The score is a combination of instructions to performers following flashing events, and conventional notation. At the same time, Jane Grant and Tim Hodgson developed the Neu-

rogranular Sampler concentrating on large grain duration, high firing density and the creation of a working interface for the instrument. The instrument formed the basis of Jane Grant's video/ audio work *Threshold* [12] which merges the spike timings with the sound of voice and breath.

8. Conclusion and Future Work

At the time of writing, one month remains of The Fragmented Orchestra's 10-week tenure at FACT. The public response has been overwhelmingly one of curious engagement: much of the the broadsheet coverage [9] has tended towards the explanatory rather than the critical, outlining the concepts, context and architecture behind the system. However, as is recognized by this article, the experience of the installation itself is actually wholly intuitive and uncomplicated; without any technical interface as such, the visitor is left simply to wander between the 24 ceiling-hung speakers, configuring their own unique experience of the installation through their location.

Timing is also critical: it quickly becomes apparent that, as David Stubbs comments in *The Wire* [10], the sonic signature of the UK – and thus the installation – vary wildly according to the time of day, with inevitable early-morning lulls, lunchtime buzz and chatter, and a greater number of intentional performances towards the evening. Though this does mean that the qualities of the experience will vary wildly according to the sound events at each of the sites, this is also one of its most rewarding features. As Stubbs comments [10]:

So fitful and sporadic are these sounds that it's left to chance just how well you'll be rewarded from your visit here - but there is a strange, radio ham's delight when, for instance, the sounds of a kid's party in a London gallery break down the wires, unspoiled even in this era of media supersaturation.

In order to promote the performative aspects of the project, the team are organising 24 hours of performance at the 24 sites on 21st February 2009 in an event called *24 Fragments*. All performances (which include music, lectures on neuroscience and distributed performances) will be relayed through the architecture of The Fragmented Orchestra and online at our website.

References

- [1] J. Prophet and M. d'Inverno. "Transdisciplinary research in CELL," in *Aesthetic Computing*, Paul Fishwick, ed. Massachusetts: MIT Press, 2005.
- [2] Steinheider, Brigitte and George Legrady. "Interdisciplinary Collaboration in Media Arts: A Psychological Perspective on the Production Process," *Leonardo* vol. 37 no. 4, pp 315-321, 2004.
- [3] D. Jones et al. "The Fragmented Orchestra website", 2008. Web: <http://www.thefragmentedorchestra.com/>
- [4] Izhikevich E.M., Gally J.A. and Edelman G.M. "Spike-timing dynamics of neuronal groups," *Cerebral Cortex* no. 14, pp. 933-944, 2004.
- [5] Gerstner W. and Kistler W. *Spiking Neuron Models: Single Neurons, populations and plasticity*. Cambridge: Cambridge University Press, 2002.
- [6] Miranda, E. R. and Matthias, J.R. "Granular Sampling using a Pulse-Coupled Network of Spiking Neurons," *Proceedings of EvoWorkshops 2005*, Lecture Notes in Computer Science 3449, pp. 539-544. Berlin: Springer-Verlag, 2005.
- [7] John Matthias and Nick Ryan. *Cortical Songs* (CD), London: Nonclassical Records, 2008.
- [8] Hodgkin, A. L. and Huxley, A. F. "A quantitative description of ion currents and its applications to conduction and excitation in nerve membranes," *J. Physiology* no. 117, pp 500-544, 1952.
- [9] A. Hickling, "Noises off," *The Guardian*, 19 December 2008. Available: <http://www.guardian.co.uk/music/2008/dec/19/fragmented-orchestra-fact-liverpool>
- [10] D. Stubbs, "Dingdong," *The Wire*, February 2009.
- [11] D. Jones. "AtomSwarm: A Framework for Swarm Improvisation," *Proceedings of EvoWorkshops 2008*, Lecture Notes in Computer Science 4974, pp. 423-432. Berlin: Springer-Verlag, 2008.
- [12] Jane Grant *Threshold* (Sound and Video) Artsway Gallery UK