

Molecular Music: repurposing a mixed quantum-classical atomic dynamics model as an audiovisual instrument

Prof. J. I. Hyde, BA(Hons), DipHE, PhD

Department of Music, Bath Spa University, Bath, UK

www.josephhyde.co.uk

e-mail: j.hyde@bathspa.ac.uk

Dr T. Mitchell, BSc (Hons), PGCert HE, PhD

Department of Computer Science and Creative Technologies, University of Western England, Bristol, UK.

Dr D.R. Glowacki, MA, PhD, FRSC, URF

School of Chemistry, University of Bristol, UK

Department of Computer Science, University of Bristol, UK

PULSE Institute and Dept of Chemistry, Stanford University, USA

SLAC National Accelerator Laboratory, Menlo Park, USA

Abstract

Molecular Music is an offshoot of a long-term collaborative project, the multi-award winning danceroom Spectroscopy (dS)[1]. dS has its roots in computational molecular dynamics, and was originally conceived to serve as an interactive platform for molecular simulation [2,3,4]. It offers a multisensory immersive experience based on cutting-edge science facilitating an understanding of the principles of our microscopic world through direct experience rather than traditional academic learning. It consists of a system of particles, simulated according to strict scientific principles; represented both visually and sonically, which can be interacted with through human movement. The project consists of a public installation, and also a contemporary dance piece, Hidden Fields, which is performed using the system. Hyde's contribution to the project consists of the sonification (interactive systems and sound design) for the installation, and the composition of an interactive score for the dance piece.

Molecular Music is intended to facilitate further exploration of the audiovisual relationships at play in dS and Hidden Fields and to explore more deeply how to sonify vibrations on a micro scale (where sound does not, as such, exist). We have built some highly developed algorithms based on FFT analysis of molecular vibration data outside the range of human hearing to yield subharmonics on which sonic material can be based. We also have in place a sophisticated system whereby sound can control the particle system and the particle system can in turn control the sound. We are exploring how this combination can be used to make a

novel kind of feedback loop, and a network of non-trivial audiovisual relationships whereby the influence of sound on image and vice versa is mediated via the medium of an advanced mixed quantum-classical model. Using these tools we can use dS as a highly evolved 'visual music' instrument.

The performance consists of a solo audiovisual performance of around 15 minutes duration. The paper outlines the algorithms at the heart of the dS system and their broader implications for Sci/Art visualisation/sonification and understanding, before moving on to examine how these algorithms have been adapted as an audiovisual instrument. The history of the project, including installations, dance performances and music-based collaborations, will be examined followed by a look to the future – in particular the development of dS as a large-scale permanent exhibit for ZKM in Karlsruhe to open in 2015.

1. The danceroom Spectroscopy project – background

danceroom Spectroscopy is an indirect outcome of the scientific research of Dr. David Glowacki. Glowacki works in the field of Computational Chemistry, using complex algorithms to model the behaviours of matter and energy (a detailed exposition of the algorithmic basis of dS is outlined in section 3 below). In 2011, with programmer Phill Tew, he began to explore the idea of a visualisation system which would afford non-specialists an insight into the principles at work in Glowacki's research, and an instinctive understanding of the inner workings of our world on a nano scale.

In its initial form, c. mid 2010, the project consisted of a simple particle system, but one where the movement of the particles is governed by the algorithms at work in Glowacki's research. These particles inhabit what Glowacki and others in his field refer to as an 'energy landscape'. One might simply think of this as a contoured landscape in which 'higher' areas (hills, mountains, ridges) are 'hotter' (ie, of higher energy) and 'lower' areas (valleys, dips, craters) are 'colder' and have lower energy.

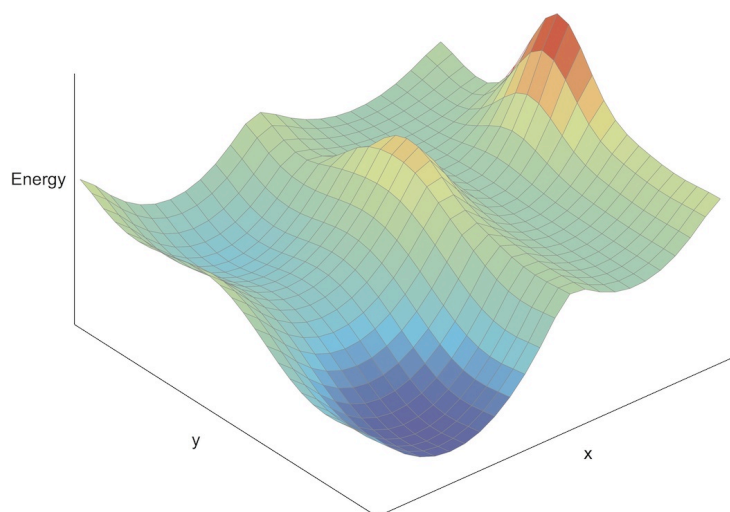


Figure 1. A representation of an energy 'landscape'

A crucial feature of the energy landscape at play in dS is that someone viewing the particle system can actually become part of this energy landscape. This is achieved in a relatively straightforward manner by interpreting their body as a peak, or a trough (both are possible with the system) in the energy landscape. Various interactive technologies have been explored in order to make this possible, but early on in the project the Microsoft Kinect sensor was adopted as a cheap and relatively simple solution.

The Kinect, introduced in November 2010, was released as a games controller for the Xbox 360 games console, but was rapidly 'hacked' and adopted for many broader purposes, including a number of arts projects. The controller consists of a standard-definition colour camera (broadly equivalent to a webcam in terms of functionality and quality), an infra-red camera, an infra-red laser projector (strictly speaking, both are near-infrared) and a microphone array. The infra-red projector and camera, working in tandem, constitute a so-called "structured light" system which allows the device to yield a depth value for every pixel of the camera image. This "depth map" is produced by comparing the pattern of dots projected with the pattern picked up by the camera, and extrapolating depth values from distortions detected in this pattern. For most applications, the purpose of this depth-map is simply to separate foreground from background, and specifically to separate human forms from what – in the domestic setting envisioned for this device – might be quite a cluttered and chaotic environment. dS uses this background separation, but also makes full use of the depth map to allow the shape of human bodies, captured in real time, to become contours in the energy landscape.

Early versions of dS were essentially manifested as small-scale public installations for demonstration purposes, using a single Kinect. However, these rapidly garnered considerable interest and in light of this the scope of the installation was expanded. The graphics were improved, more Kinects were used to allow the system to work in larger spaces (an array of up to 10 of the devices can be used with the latest version) and – crucially for the purposes of this paper, a sonic component of the project began to be developed. The latter will be discussed in detail below.

2. Hidden Fields

By 2011, the installation had grown into a highly successful and popular Sci-Art installation sufficiently flexible to be installed in a wide variety of spaces, including a dome environment with full 360-degree capture and projection (see figures 1 and 2). At this point a large team of collaborators had gathered around the project, including Professor Joseph Hyde and Dr Tom Mitchell. A decision was made at this point to further explore the aesthetic possibilities of the system using a small number of 'expert users'. Trained dancers seemed the obvious choice for such a role. Applicants from the contemporary dance community were invited to apply to take place in a specialist workshop using the system at the Arnolfini in Bristol in June 2011. Following the success of this a decision was made to

develop a fully-fledged dance work, known as Hidden Fields (see figure 3). This went through a number of iterations throughout 2012, with development supported by a grant from the Arts Council of England, and a number of performances in the UK and Europe.



Figure 1 – the Igloo 360 dome being installed for the dS festival, The Passenger Shed, Bristol, October 2012



Figure 2 – the dS installation in dome formation



Figure 2 - Hidden Fields performance in Seeing Sound [5], November 2013

Hidden Fields offered a tremendous opportunity to develop the sonic and musical dimensions of the project. During this period, Hyde and Mitchell, broadly operating in the roles of composer and music technologist, evolved the broad palette of sounds and techniques outlined below. The controlled environment of a dance performance and an extensive rehearsal process incorporating dancers (as expert users) allowed for an exploration of extremes. Whilst a public installation needs in some respects to 'play it safe', to allow for a wide range in terms of the numbers of users and the ways in which they might interact with the system, Hidden Fields allowed us to decide what we wanted the system to do, and to evolve ways of achieving this end. In this way, the aesthetics of the project were allowed to evolve, visually and sonically.

As the piece evolved towards its current duration of around an hour, it was felt necessary to firm up a quasi-narrative structure for the piece. This was largely driven by the music, and we felt the need to have certain points in the structure where the sound/music was not fully driven by the particle system (with the chaotic behaviour that paradigm produces) but rather allowed to operate according to a more traditionally musical/compositional logic. Rather than break the connection between the particle system and the sound at these junctures, we evolved a set of techniques whereby the sound could control the particle system as opposed to vice versa. This opened up a set of possibilities only explored in embryonic form in Hidden Fields, and was arguably the origin of what would become Molecular Music. It allowed the possibility of a complex feedback loop where the particle system and sound operate upon each other. We only scratched

the surface of what such a system makes possible within the context of Hidden Fields and this led us to actively seek another avenue within which to do so.

3. Scientific Underpinnings

As stated above, the Kinect sensor (or array) yields a depth map whereby for every pixel in an x/y array, a value is given for z (depth). The plot in figure 5 shows a human form incorporating data derived from this depth map - here the intensity of the colours is linked to the magnitude of the local gradient vector on the image. The manner in which it is plotted suggests analogy with the concept of an energy landscape, which has become a fundamental idea guiding how chemists and physicists think about both kinetics and dynamics in a range of chemical systems, from small molecules to complex materials and biochemical systems [6, 7]. An energy landscape is effectively a topological map of a system's potential energy, V , at a range of different configurations. Within any localized region of the energy landscape, the gradient of the energy, $dV/d\mathbf{q}$, relates the topology of the energy landscape to the classical forces felt by a particular molecular configuration. dS interprets people's movements as perturbations on a virtual energy landscape.



Figure 5 - Force topology map of the human form. Gray indicates a gradient of zero. The intensity of each color is related to the magnitude of the local force vector on the depth image. Color choice has been selected to effectively illustrate depth

In its present form, dS carries out an MD simulation involving N atoms, each of which may move in a virtual coordinate system defined by Cartesian x , y , and z directions. Hamilton's equations of motion, commonly used to discuss the dynamics of molecular systems in both classical and quantum frameworks, provide a useful vantage point for describing how the system works. They are as follows:

$$\begin{aligned} d\mathbf{p} / dt &= -dH / d\mathbf{q} \\ d\mathbf{q} / dt &= dH / d\mathbf{p} \end{aligned} \quad (1)$$

where \mathbf{p} and \mathbf{q} are the momentum and coordinate vectors of each atom in the ensemble, and H is the so-called Hamiltonian function describing the total system energy - i.e.:

$$H = \sum_{i=1}^N \frac{m_i v_i^2}{2} + V \quad (2)$$

where i is an index that runs over a collection of N total atoms, m is the mass of an atom, and v is its velocity. The first term in Eq (2) describes the total kinetic energy of the system while the second, V , describes the total potential energy. Within dS, there are two different contributors to V :

$$V = V_{\text{int}} + V_{\text{ext}} \quad (3)$$

where the total potential energy, V , is calculated as the sum of two terms, V_{int} and V_{ext} , which correspond to the potential energy owing to internal and external fields, respectively. Like many MD programs, the most expensive loop in dS is associated with calculating V_{int} , and involves summing over all possible pairwise interactions:

$$V_{\text{int}} = \sum_{i=1}^N \sum_{j=i+1}^N V(r_{ij}) \quad (4)$$

where r_{ij} is the distance between atoms i and j . During initial prototyping and benchmarking of the dS system, V_{int} included only non-bonded Lennard-Jones type interactions with parameters derived from electronic structure calculations [8]. However, as discussed further below, we have recently implemented a set of fast C# wrappers which allow dS to call the GPU-accelerated OpenMM program whenever a force evaluation is required. OpenMM allows for a wide range of force interactions, including bonds, angles, torsions, non-bonded Lennard Jones interactions and electrostatic interactions [12].

The V_{ext} term in Eq (3) is calculated as a sum over the difference between a raw depth matrix at time t , $V_{\text{ext}}(x_i, y_i, t)$, and an average background depth image, $\langle V_{\text{ext}}(x_i, y_i, 0) \rangle$ as follows (angled brackets indicate an average):

$$V_{\text{ext}} = C_i \sum_{i=1}^N [V_{\text{ext}}(x_i, y_i, t) - \langle V_{\text{ext}}(x_i, y_i, 0) \rangle] \quad (5)$$

where the term in square brackets represents the potential energy that an atom ‘feels’ as a consequence of people’s motion, and C_i is a variable scaling constant applied to a specific atom. Interactive control over C_i allows the user to determine how strongly any given atom ‘feels’ forces from the users’ fields, and whether a person’s field is ‘attractive’ or ‘repulsive’. Eq (6) is responsible for coupling human motion to the atomic dynamics, allowing humans to sculpt the potential energy landscape felt by the atomic ensemble, and thereby chaperone the system dynamics.

In Hamiltonian mechanics, the energy function, H , should remain constant for any closed dynamical system, in line with the conservation of energy required by the first law of thermodynamics [7]. However, the Eq (2) Hamiltonian is not subject to this constraint because of the V_{ext} term, which effectively makes the system open rather

than closed. Fluctuations in the depth data arise as a consequence of noise in the depth images, or people's motion within the space mapped by the depth sensors. Both of these effects effectively result in fluctuations of the total system energy, introducing significant instabilities into the Velocity Verlet [10] scheme used to propagate the time-dependent system dynamics in Eq (1). To address this, and avoid the crashes associated with such instabilities, we have implemented a modified velocity rescaling Berendsen thermostat, in which the instantaneous system temperature T_t approaches some desired temperature T_0 with a first order rate

$$\frac{dT_t}{dt} = \frac{1}{\tau} \cdot (T_0 - T_t) \quad (6)$$

that depends on a user-specified rate coefficient ($1/\tau$) and how far the system is from T_0 . Rearranging (6) gives an expression for the temperature change dT_t over some time step dt :

$$dT_t = \frac{dt}{\tau} \cdot (T_0 - T_t) \quad (7)$$

where τ is a first order time constant, and

$$T_t = \frac{1}{d \cdot N \cdot k_B} \sum_{i=1}^N m_i v_i^2 \quad (8)$$

with d the number of dimensions in which each atom can move (three), N the number of atoms in the simulation, and k_B the Boltzmann constant. The velocity rescaling constant λ is determined via definition of $T(\lambda)$, which is the temperature that results when all the atomic velocities are scaled by λ , i.e.:

$$T(\lambda) = \frac{1}{d \cdot N \cdot k_B} \left[\sum_{i=1}^N m_i (\lambda v_i)^2 \right] = \lambda^2 T_t \quad (9)$$

We evaluate λ by specifying that $dT_t = T(\lambda) - T_t$, and substituting Eq (7) and (9) to give:

$$\frac{dt}{\tau} \cdot (T_0 - T_t) = \lambda^2 T_t - T_t \quad (10)$$

which may be solved to yield

$$\lambda = \sqrt{1 + \frac{dt}{\tau} \left[\frac{T_0}{T_t} - 1 \right]} \quad (11)$$

Prior to determining the value of T_t required for calculating that atomic velocity scale factor λ , there is an added stability measure: we loop over the atomic velocities to ensure that none of the atoms within the simulation have a velocity more than two

standard deviations larger than the average atomic velocity. We have found the procedure outlined above gives a good compromise between computational efficiency, interactive fluidity, and system stability. Moreover, it is extremely robust to numerical instabilities that can arise when user motion suddenly ‘injects’ energy into the system Hamiltonian.

4. Technology

The primary dS software consists of a bespoke application coded in C#. Later versions of the software are highly CPU-optimised, and make full use of top-end NVIDIA graphics cards – the primary system currently in use contains a dual-card system with over 4000 cores, all of which are co-opted by the software. Thanks to the work of a team of programmers specialising in code for GPUs, these cores are used not only to produce the graphics for the system, although at up to 6xHD resolution (11,520 x 1080 pixels – see figure 6) the rendering is a non-trivial task.




Figure 6 – 11,520 x 1080 pixel renders captured directly from dS software

The GPU cores are also used to resolve the equations outlined above – a unique aspect of the project which has attracted considerable interest from the scientific community, since it essentially allows a domestic PC (albeit a very high end one) to function as a massively parallelized supercomputer for the solving of many simultaneous quantum dynamics algorithms in real time. For the purposes of dS, it has allowed us to massively increase the particle count.

The audio component of dS runs on a separate machine, and is built using Cycling74's Max environment [11]. For the purposes of Hidden Fields and subsequently Molecular Music, the Max software has been devolved into a number of Max for Live [9] plugins. The advantage of this approach is that it allows Ableton Live to function as a timeline for durational structures – whilst the micro-level elements of both sound and image will be aleatoric and interactive, the overall behaviour of the system can be controlled with the sophistication and reliability that a modern Digital Audio Workstation (and in this case, one specifically designed for live performance) affords.

Communication between the two machines is achieved using the Open Sound Control (OSC) protocol, developed by Martin Freed and his team at CNMAT, University of California, Berkeley [13]. This protocol offers many advantages for us in this context – primarily by virtue of its flexibility and speed. We have been

able to develop a system where what multiple streams of information are sent in both directions – these can be massively parallel (for example, collision data from every single particle in the system) and/or extremely fast – we are able to use data streams at close to audio rates, which makes the system incredibly responsive and opens up many creative avenues.

Although the OSC protocol operates essentially as a network, allowing data to travel in any direction between multiple machines, dS – at least as used in Hidden Fields and Molecular Music – operates in a configuration where the PC running the main dS software (quantum simulation, particle display, depth map acquisition etc.) is essentially ‘slaved’ to or controlled by the audio machine. This is mentioned here because it is one of the factors that makes the system so suitable as a compositional or audiovisual performance tool – this is discussed in more depth below.

5. Sonification – general principles

We will discuss in turn the two primary categories of interaction involved in the system: firstly interaction where data is being sent from the primary dS machine to the audio machine and secondly where data is sent in the opposite direction. In reality both are operating continuously and simultaneously, as will be discussed below.

The simplest type of interactivity, and the first to be developed, is collision detection. Data is produced when particles collide, at which point the system yields the coordinates at which the collision takes place, and the speed at which the particles collide. The capabilities offered by the type of interaction will vary according to the physical properties of the system being simulated (this is something of a universal principle in terms of dS, and is something we discuss further below). Where particles and energy fields are ‘attractive’, ie where simulated physical forces will draw them together (using the energy landscape metaphor, this would be equivalent to a ‘trough’), particles will tend to be drawn together to form clusters, and the collision data – when sonified – will tend to yield recognisable patterns; not entirely regular but nonetheless with a clear contour. These can function well as stochastic rhythms or melodies.

Where simple particle collision data is combined with a model whereby particles and energy fields repulse each other (modelling ‘peaks’ on an energy landscape), a different set of possibilities is afforded. With no external stimuli the system in this state will become entirely chaotic, but with an external source, the particles will tend to be driven by boundary conditions in the source (edges between areas of different hue or luminance, for example) and to form wave-like structures at these points. Within the sonification, this will tend to yield granular-type textures where the sounds associated with individual particles get lost perceptually, but overall contours of

particle waves are perceptible and structurally useful. The fact that this type of sonification only becomes perceptually interesting with some kind of external stimuli is why a video source is still usually used with Molecular Music performances as a 'substitute' of sorts for the Kinect-derived depth map used in the original dS installation. The content of the video source is almost irrelevant – it can simply serve to provide some kind of irregularity or granularity in an energy landscape that would otherwise be entirely uniform.

Once a large number of particles are introduced it can be very hard to produce structures and sounds which are meaningful using collision data – as might be imagined, if tens of thousands of particles and the collisions between them are sonified the result will tend to approach noise. At this point there are two other models of interactivity that can be employed. The first of these has become known as 'group data'. This data will generally only be produced, or at least meaningful, under certain conditions. There will need to be a relatively large number of particles, and they will need to be 'attractive' (ie attracted to energy fields). In this set of circumstances, the particles will tend to 'swarm', to form into large groups. (In the installation or dance performance, these groups will tend to correspond to individual users). The system has the capability to treat such groups as 'super-particles', and the system can provide corresponding data to that which it provides for individual particles - spatial coordinates and speed data. This allows the production of relatively simple, or at least perceptually 'digestible' musical patterns and contours using very large numbers of particles. However, it is dependent on a fairly specific set of conditions in terms of the physical simulation – without these the groups will not even form.

Where large numbers of particles are used in more chaotic formations, without the grouping mentioned above, a third type of data becomes particularly useful. This data involves a Fast Fourier Transform (FFT) analysis of the vibrational energy of the particle system (the frequencies yielded by the transform function are sub-audio, below 1 Hz, but are effectively transposed up into the audio domain). This is of interest because it will tend to highlight overall properties of the particle system, in particular any kind of coherence in the movement of the particles in the system.

Such coherent movement will result in measurable peaks in the FFT data, which can be made to yield perceptible sonic feedback when transposed up into audible frequency space. In the case of installation or dance performance, these peaks will be produced when a significant number of the users in the space move in the same direction, or in the same manner. In the case of Molecular Music, in the absence of such stimuli, the only way to produce such coherent motion is by using sound itself to stimulate the particles to move in an ordered fashion – this forms the basis of a kind of feedback loop which has become core to the project.

At the same time as these methods are provided whereby the particle system can control the sound, as of latter versions of dS, the sound can also control the particle system – we have explored methods whereby amplitude, frequency and more complex FFT-derived data can be used to control all aspects of the simulation. What is particularly interesting is that the sound is controlling physical rather than graphical properties, so that (for example) a mapping of sound intensity onto temperature will manifest itself as complex and multifarious (and sometimes unpredictable) changes to the visualisation.

In addition to mapping sound-derived data onto physical properties, we have also built in functionality whereby the sound machine and software can control the dS machine directly through OSC. This is largely used to achieve long timescale changes using the timeline functionality of Ableton Live (and Max for Live). It allows us to shape the system into time-based structures – essentially, to ‘compose’ it. This has proved invaluable in the Hidden Fields dance project and Molecular Music.

Of course, all these types of sonification and visualisation (in this context these terms, particularly the latter, could be seen as something of a simplification) are used in various permutations, often all together. Recent versions of the dS system allow the combination of ‘attractive’ and ‘repulsive’ behaviours, where some types of particles exhibit the former pattern of behaviour and some the latter. The permutations between these types of behaviour and the three types of data offer a plethora of possibilities, and make the dS system uniquely flexible in terms of sonification.

6. Molecular Music – Introduction

Molecular Music started life through the agency of a specific opportunity. The dS team were offered the chance to perform with violin virtuoso Nicola Benedetti and her string trio, as part of the first-ever Bristol Proms season at the Bristol Old Vic in August 2013 (figure 7). Given limited rehearsal time for this performance, we used something close to the existing dS setup, but with additional weight given to sound-particle system interactivity. We used the system with a multichannel interface for the first time, and used separate feeds from violin, cello and piano to drive individual elements of the simulation.

This performance was considered sufficiently successful that we were invited back to participate in the 2014 Bristol Proms. Tom Morris, the festival’s Director, agreed that it would be interesting to explore a more flexible and collaborative framework in this second iteration, and to have more time to develop a unique working methodology. In this instance we worked with the Charles Hazlewood All-Star Collective, a loose confederation of musicians mostly based in Bristol and the South West. The crucial difference between this performance and the one that

preceded it was that whereas the Benedetti trio performed from the Classical music repertoire, Hazlewood and his ensemble worked up an interpretation of Terry Riley's 'Rainbow in Curved Air'. This work has much in common with Riley's work 'In C' in structure, but actually has no score – our performance was based on an interpretation of the 1969 album (recorded by the composer) and the 'Rainbow in Curved Air Calligraphies' – abstract representations that Riley released more as interpretations of the music than as performance instructions.



Figure 7 - The Nicola Benedetti trio at the Bristol Proms with dS visualisations



Figure 8 – Bristol Proms performance with Charles Hazlewood All Star Collective

This time we were able to enter a more collaborative process where improvising musicians were able to respond to the visual stimuli yielded by dS as much as dS was in turn responding to their playing. Although this was some way short of a feedback loop (dS only providing the visual dimension to the performance rather than driving any sonic elements) it provided the closest analogy to some of the principles that would become core to the Molecular Music endeavour. It also allowed us to develop considerably more sophisticated sonification algorithms – in this instance we had separate audio feeds from 8 musicians, as well as several channels of MIDI, and a research and development period at Bath Spa University with some of the musicians allowed us both to evolve these techniques and to give the musicians time to learn how to ‘play’ the system to an extent.

One of the most important departures in the *Rainbow* performance was the decision to jettison the Kinect array. This put the emphasis firmly on sonification – although we did actually use a standard camera feed (a moving camera trained on the musicians) to provide some texture to the simulation in the manner outlined above, this was thoroughly abstracted and had no depth map – the energy landscape employed by the simulation was therefore entirely driven by sound rather than image.

7. Molecular Music – Aesthetic Considerations

Molecular Music essentially reimagines danceroom Spectroscopy as an audiovisual instrument. We see it as sited broadly as within the field of Visual Music, defined by Hyde thus: *Visual Music involves the artistic expression of musical ideas or material through ocular media* [15]. Although this term is generally applied to fixed media artworks, whether they be paintings, cinema or video art, many interactive and/or performative examples can be found, dating as far back as Oskar Fischinger’s *Lumigraph* in 1950 (and of course earlier ‘light organ’ instruments, and far more prevalent in more recent times). These tend to be interactive – that is to say user controlled, or to be based primarily on visualisation (visual elements driven by sound) or sonification (vice versa). Although examples can be found that combine more than one of these modes, Molecular Music is unusual in that it combines all three; even more so in that the interaction between these modes is via the medium of a quantum molecular simulation.

This factor is crucial in determining the characteristics of Molecular Music as a visual music instrument. In the design of any interactive system it can be hard to find the sweet spot between ‘mickey mousing’, where mappings and relationships are highly perceptible but too facile to be of lasting interest, and complex interactions where these relationships may be too indirect to be easily perceived. Using the simulation as a medium allows for a certain amount of uncertainty to be introduced which gives a pleasingly organic characteristic to modal interactions. We might use the metaphor of a vactrol, a component popular in modern analogue synthesizer designs. A vactrol contains an LED and a light dependent resistor – a varying voltage (or sound) at the input controlling the brightness of the LED will produce a variable resistance from the resistor, usually converted back

to a voltage variation via a voltage divider circuit. The indirect nature of the connection between input and output (actually, in electrical terms, they are completely disconnected) results in a certain amount of variability that is highly prized as imparting a kind of musicality.

This is a function of principle that the system is predictable but also chaotic. This is evident even in the simplest instances. Figure 9 shows two runs of the same simulation, and the most basic that dS can offer. In this instance, a single particle is released and given an initial velocity by a short burst of a 1000 Hz sine wave (producing a short 'spike' in the simulation temperature). It loops back on itself by virtue of its own mass and, as can be seen, the basic form in each case is the same but the exact details of the form differ. Every time this simulation is run the results are recognisable but unique. This is perhaps not surprising from a scientific perspective, but from an artistic viewpoint it has great potential in injecting a 'musical' variation into the interaction between sound and image.




Figure 9: two instances of a one-particle simulation

As the complexity of the system increases, so do the creative possibilities of this chaotic behaviour. Figure 10 shows a very similar situation to that seen in figure 9 above, - again a short sine tone burst is controlling the temperature of the simulation, but in this case with many more particles of several types.

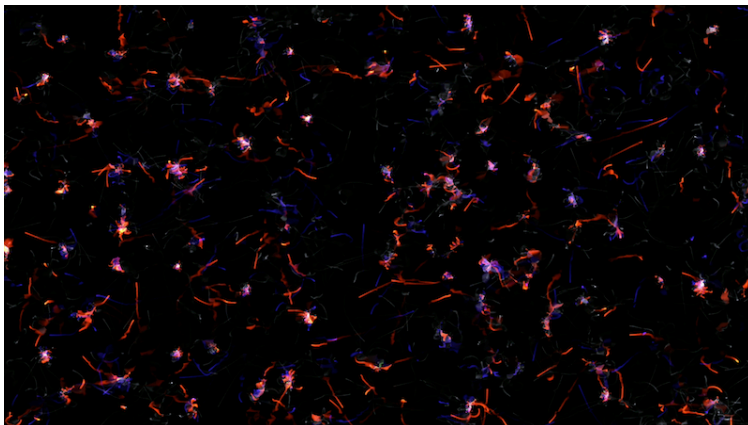




Figure 10: a similar simulation with multiple particles of different types

This is particularly the case where a feedback loop is established. The basic methodology of such a feedback loop is to set up conditions in which the particles control the sound and the sound controls the properties of the simulation (note that this is not directly equivalent to sound controlling image and image controlling sound). This amplifies the chaotic qualities of the system. The results are far less predictable, but are nonetheless repeatable. With exploration aesthetically pleasing results can be found. Some examples are shown here:



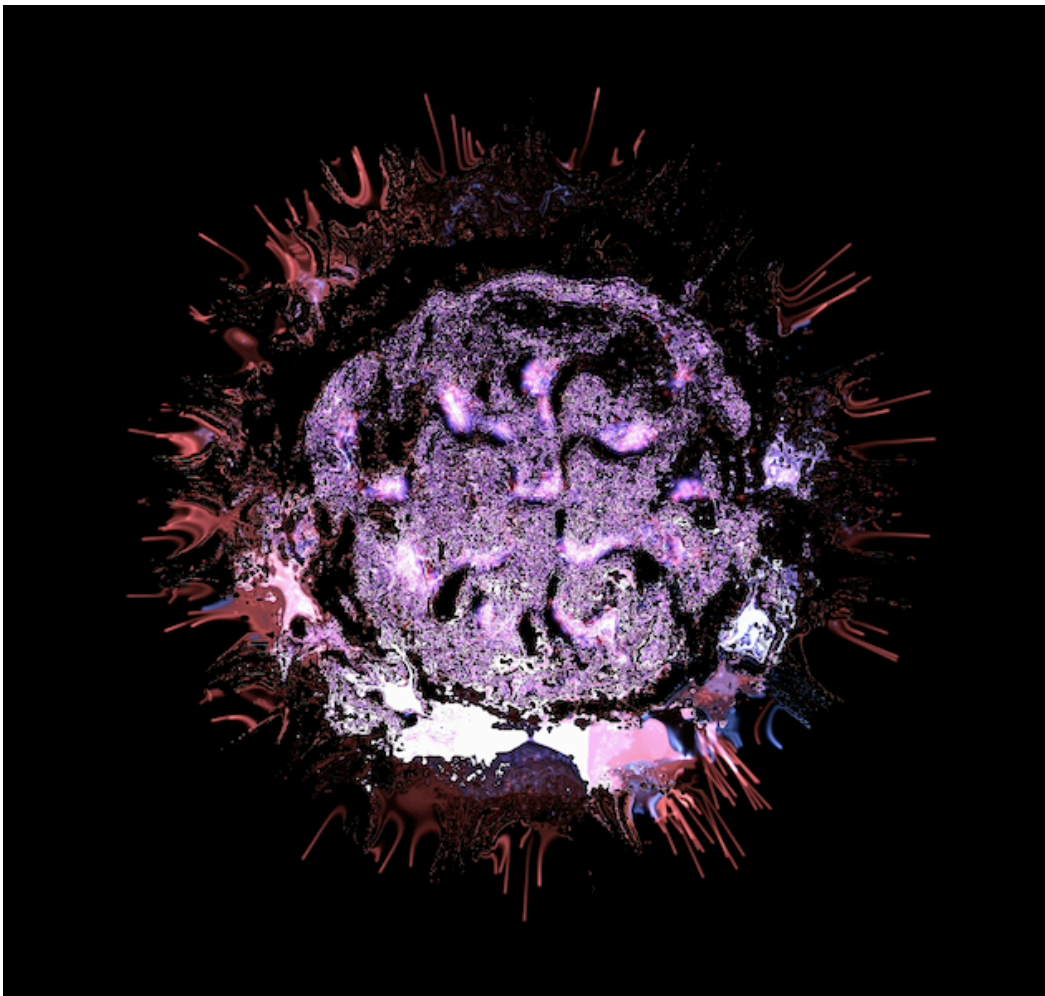



Figure 11 – Complex forms made through feedback via the dS simulation

8. Next Steps

The danceroom Spectroscopy project and its 'satellites' Hidden Fields and Molecular Music continue to enjoy great success and exposure. The scientific basis of the simulation is being continuously developed by Dr Glowacki and his colleagues at Stanford University. The dS system described here, and used for Molecular Music performances to date, only simulates single atoms, whereas versions of the software have since been produced to simulate more complex molecules, even protein strings (see figure 12 below)




Figure 12 – Interactive protein folding using a dS-based simulation

Work is also being carried out to allow more complex energy 'landscapes' to be explored using the system. A recent experiment (shown in figure 13) allows users to literally sculpt such a landscape using a sandbox, and then to see how particles of matter (literally projected onto the sand landscape) would behave in the landscape created. Neither of these recent developments has incorporated a sonic element as yet, and the possibilities for sonification seem very promising.




Figure 13 – sculpting an energy landscape in sand

A great opportunity will be offered by the building of dS as a permanent installation, at ZKM in Karlsruhe, Germany. There are many challenges involved in such an undertaking – reliability and sustainability to name but two. Amongst the more creative challenges will be the sonic aspects of the installation. As an installation, dS has to date usually involved a single ‘state’ (model of sound-simulation interaction). This has been seen as desirable in a situation where many parameters are outside of our control – the most influential perhaps being the number of people present in the space. The state that we have is capable of surprising variety, but nonetheless in an installation that may be in place for years rather than days we feel the need to allow for more long term evolution of the sound (amongst other things). This will give us the creative impetus (and space and time) to take the sonic aspects of the project to the next level.

References:

- [1] <http://danceroom-spec.com>
- [2] D. R. Glowacki et al., “A GPU-accelerated immersive audiovisual framework for interactive molecular dynamics using consumer depth sensors,” *Faraday Discussion* 169, 2014, p 63 - 89 (open access)
- [3] D. R. Glowacki et al., “Sculpting molecular dynamics in real-time using human energy fields,” in *Molecular Aesthetics*, ISBN: 9780262018784 (MIT Press), ed. Peter Weibel and Ljiljana Fruk, Sept 2013
- [4] T. Mitchell et al., “danceroom Spectroscopy: at the frontiers of physics, performance, interactive art, and technology,” *Leonardo*, doi:10.1162/LEON_a_00924
- [5] <http://www.seeingsound.co.uk>
- [6] D. R. Glowacki, J. N. Harvey and A. J. Mulholland, *Nat. Chem.*, 2012, 4, 169-176.
- [7] D. J. Wales, *Philos. Trans. R. Soc. A-Math. Phys. Eng. Sci.*, 2012, 370, 2877-2899.
- [8] N. L. Allinger, Y. H. Yuh and J. H. Lii, *Journal of the American Chemical Society*, 1989, 111, 8551-8566.
- [9] P. Eastman, M. S. Friedrichs, J. D. Chodera, R. J. Radmer, C. M. Bruns, J. P. Ku, K. A. Beauchamp, T. J. Lane, L. P. Wang, D. Shukla, T. Tye, M. Houston, T. Stich, C. Klein, M. R. Shirts and V. S. Pande, *Journal of Chemical Theory and Computation*, 2013, 9, 461-469.
- [10] D. Frenkel and B. Smit, *Understanding molecular simulation: from algorithms to applications*, Academic press, 2001.

[11] <http://www.cycling74.com>

[12] <https://www.ableton.com/en/live/max-for-live/>

[13] <http://opensoundcontrol.org/introduction-osc>

[14] Hyde, J., (2012) 'Musique Concrète Thinking in Visual Music Practice: Audiovisual Silence and Noise, Reduced Listening and Visual Suspension'. *Organised Sound*, 17/2: pp. 170-178.

[15] Keefer, C. and Guldemon, J. eds. (2012) *Oskar Fischinger 1900 – 1967: Experiments in Cinematic Abstraction*. Pp. 193 – 200.