Towards Molecular Musical Instruments: Interactive Sonifications of 17-Alanine, Graphene and Carbon Nanotubes

Thomas J. Mitchell
CT Lab, University of the West of England, Bristol, UK
tom.mitchell@uwe.ac.uk

Alex J. Jones
Intangible Realities Laboratory
University of Bristol, UK
alex.j.jones@bristol.ac.uk

Michael B. O'Connor
Intangible Realities Laboratory
University of Bristol, UK
mikeoconnor0308@gmail.com

Mark D. Wonnacott
Intangible Realities Laboratory
University of Bristol, UK
m.wonnacott@bristol.ac.uk

David R. Glowacki
Intangible Realities Laboratory
University of Bristol, UK
glowacki@bristol.ac.uk

Joseph Hyde
School of Music and Performing Arts
Bath Spa University, UK
j.hyde@bathspa.ac.uk

ABSTRACT

Scientists increasingly rely on computational models of atoms and molecules to observe, understand and make predictions about the microscopic world. Atoms and molecules are in constant motion, with vibrations and structural fluctuations occurring at very short time-scales and corresponding length-scales. But can these microscopic oscillations be converted into sound? And, what would they sound like? In this paper we present our initial steps towards a generalised approach for sonifying data produced by a real-time molecular dynamics simulation. The approach uses scanned synthesis to translate real-time geometric simulation data into audio. The process is embedded within a stand alone application as well as a variety of audio plugin formats to enable the process to be used as an audio synthesis method for music making. We review the relevant background literature before providing an overview of our system. Simulations of three molecules are then considered: 17-alanine, graphene and a carbon nanotube. Four examples are then provided demonstrating how the technique maps molecular features and parameters onto the auditory character of the resulting sound. A case study is then provided in which the sonification/synthesis method is used within a musical composition.

CCS CONCEPTS

• **Human-centered computing** → **Mixed / augmented reality**; **Sound-based input / output**; Auditory feedback; Activity centered design.

KEYWORDS

virtual reality, augmented reality, sonification, game audio, spatial audio, sonic interaction design, musicology, sound art

ACM Reference Format:
Thomas J. Mitchell, Alex J. Jones, Michael B. O'Connor, Mark D. Wonnacott, David R. Glowacki, and Joseph Hyde. 2020. Towards Molecular Musical Instruments: Interactive Sonifications of 17-Alanine, Graphene and Carbon Nanotubes.

1 INTRODUCTION

Digital simulations of acoustic musical instruments are often derived from physics-based models of vibrating objects such as strings, bars, air cavities and membranes. These methods often employ computational approximations of the wave equation such as digital waveguides [8, 35] or finite difference methods [1, 37]. As with all matter, the macroscale properties of sounding objects are determined by the microscopic properties of the underlying molecular structure, which is itself characterised by perpetual dynamism and movement. Atoms and molecules are in constant motion, with vibrations and structural fluctuations occurring at very short timescales and corresponding length-scales. This spatial and temporal dynamism far exceeds the limits of human perception; however, scientists now use rigorous computational molecular simulations as a valuable complement to experiment in order to understand and make predictions about the microscopic world. The vast majority of molecular research relies predominantly on visualisation and auditory representation is, with a few exceptions (see Section 1.1), conspicuous by its absence. We are developing generalised molecular sonification methods that are designed to convey the dynamics of the underlying simulation data without neglecting aesthetics and musical form. In this paper we present a molecular sonification approach that uses scanned synthesis to translate real-time trajectory data into sound. In Section 1.1 and 1.2 we review the literature relating to molecular sonification along with synthesis approaches for translating geometric data into sound. In Section 2 we provide an introduction to the Narupa molecular dynamics framework before outlining a scanned synthesis approach for sonifying and interacting with simulated molecules. In Section 3 we apply this technique to sonify and build synthesisers based upon three molecular simulations: 17-alanine, graphene and a carbon nanotube. In Section 4 a set of audio and video examples are presented. Section 5 provides a case study reflecting on the use of the system within a musical context. Conclusions and areas for future work are set out in Section 6.

1.1 Molecular Sonification and Music

Molecular simulations produce large, rich and dynamic datasets that are notoriously difficult to comprehend and analyse [13]. Visualisation conventions are abundant and have evolved over time, whereas auditory representations have fewer customs to rely upon. Sonification is slowly emerging as a technique to supplement molecular dynamics visualisations, with general purpose tools emerging for widely used environments like PyMOL [17], MegMol [33] and UCSF Chimera [4]. Sonification methods have also been developed

for specific molecular applications; for example, browsing RNA structures [19] or understanding molecular conformational stability [3] and representing the structural and dynamic properties of DNA [38]. Recent efforts have also been made to translate the structural features of proteins into note sequences [11] and using these sequences to generate novel proteins and music [44].

The microscopic dynamics of atoms and molecules can never be perceived directly and therefore visual and auditory representations inevitably combine scientific intention with some form of aesthetic judgement [20]. Indeed, the aesthetics and creative merit of visual and auditory data representations have been widely acknowledged to play an important role in enhancing their efficiency, usability and expressive quality [18, 22, 30, 39]. This creative aspect of scientific representation has provided a source of inspiration for interactive and data-driven music and art [28]. For scientific applications, developers of auditory displays often use *audification* [9], a sonification technique where data streams are translated directly into audio signals. An alternative method is *parameter mapping sonification* [16], where data features are mapped to control auditory parameters. These approaches are often critiqued for their functional approach, which can neglect the musical and auditory aesthetics of the display. Consequently, sonifications can be fatiguing [42] and can sometimes lead users to disable auditory feedback entirely [36]. We intend to create a generalised sonification approach to complement existing molecular visualisations that is faithful to the data fluctuations of the underlying simulation while also attending to the aesthetics and musicality of the resulting sound. Consequently, in this paper we consider how the dynamical, geometric data that is produced by an interactive molecular dynamics simulation could be used as the basis for audio synthesis and music making.

1.2 Geometric Synthesis

A number of techniques have been proposed for translating spatial objects and physical systems into sound. For example, Wave-terrain synthesis is a technique for generating audio from a multidimensional surface (or *terrain*) [5]. The terrain defines a metaphorical landscape, which is mapped into audio using a secondary *trajectory* function. The trajectory represents a path through the landscape and sound is synthesised by traversing the path at the audio rate where the altitude at each time-step is taken as the instantaneous amplitude. Dynamic timbres can be produced by modifying the path of the trajectory through the terrain. Since wave-terrain synthesis assumes a static terrain, it is inappropriate for generating audio from the rapidly changing trajectories produced by molecular simulations. However, a closely related technique is scanned synthesis: where a dynamic model of a physical system, updating at subsonic rates, can be scanned along an arbitrary path at audio rates [41]. Scanned synthesis has previously been used to interactively sonify digital images [27], mobile sensor data [26] and geometric shapes [32]. Tubb et al. developed a system called *The Wablet* [40], a touchscreen application that uses scanned synthesis to sonify a simulated mass and spring network (or *mesh*). The mesh is visualised in real-time and updated at the video frame rate (~60 Hz), while sound is generated by scanning a path through adjacent network masses at the audio rate. The system operates using a dynamic wavetable approach, where the euclidean displacement of each mass along the scan path is converted to an amplitude value in the wavetable. The timbre and dynamics of the sound reflect the instantaneous state of the mesh, which can be manipulated in real-time via its visual representation using touch interactions.

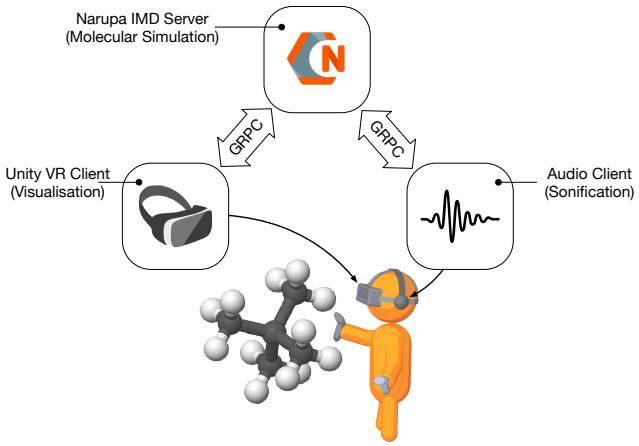

Figure 1: System block diagram showing Narupa server, VR and audio clients

2 MOLECULAR MUSICAL INSTRUMENTS

Motivated by the work above, we have developed a sonification technique that uses scanned synthesis to generate audio from a real-time molecular dynamics simulation framework called Narupa.

2.1 Narupa

Narupa [23, 24, 31] is an interactive visualisation framework with a focus on VR. The system enables researchers to observe and interact with simulated atoms and molecules that are represented as 3D dynamical objects, vibrating and moving according to rigorous real-time approximations of the physics that govern their motion. Narupa simulations are configured in the scripting language Python and can run locally or remotely on appropriate hardware, such as cloud or high performance computing clusters. The Narupa iMD VR client application allows researchers to intuitively 'reach into' the simulation in real-time and manipulate the molecular structures as the system runs [23]. As represented in Figure 1, connected client applications are continuously streamed updates from the server (e.g. atom positions, kinetic energy, etc.) and can similarly send simulation interactions and commands back to the server. Server and client communication is implemented using GRPC, a high-performance remote procedure call system supporting bidirectional streaming in multiple programming languages [15].

The simulations in Section 3 use the Narupa integration with Atomic Simulation Environment (ASE) [25] and the GPU-accelerated OpenMM framework [10] to compute atomic positions, which are streamed to connected clients for visual or auditory rendering, see [31] for further details. The dynamics are simulated in ASE under Langevin molecular dynamics [12]. For a system of N atoms, given a target temperature T and a friction coefficient γ, the Langevin equation governs the acceleration of each atom, $\vec{a} \in \mathbb{R}^{3N}$, with the following equation:

$$\mathbf{M}\vec{a} = -\vec{f} - \gamma\vec{v} + \sqrt{2\gamma k_B T}\vec{R}, \tag{1}$$

where $\vec{f} \in \mathbb{R}^{3N}$ is the vector of forces acting on each atom, $\vec{v} \in \mathbb{R}^{3N}$ is the vector of velocities of each atom, k_B is the Boltzmann constant and $\vec{R} \in \mathbb{R}^{3N}$ is a source of Gaussian random noise and $\mathbf{M} \in \mathbb{R}^{3N\times 3N}$ is the diagonal matrix of atomic masses.

The system is interactive as it is possible to apply 'external' forces to individual atoms and consequently the forces acting on simulated atoms consist of two parts: the internal forces $\vec{f}_{int}$ and the external interactive forces $\vec{f}_{ext}$:

$$\vec{f} = \vec{f}_{int} + \vec{f}_{ext}. \tag{2}$$

The internal forces, for this work, are modelled with an MM3 molecular mechanics force field in OpenMM, describing how the atoms interact with each-other via bond, angle and non-bonded terms [2].

External spring forces can be set to act on an individual atom j from any 3D position, $\vec{g}_i \in \mathbb{R}^3$ via:

$$\vec{f}^{j}_{ext} = 2m_j c\left(\vec{r}_j - \vec{g}_i\right). \tag{3}$$

where $\vec{f}^{j}_{ext} \in \mathbb{R}^3$ is the force to be applied to atom j, c is a scale factor that tunes the strength of the force interaction and m_j is the mass of the selected atom. In a VR context, this mechanism enables users to interact with the simulation by pulling the triggers on their hand controllers which exerts a force on the nearest atom. For the purposes of audio synthesis, the external forces provide a means of perturbing and injecting energy into the system, analogous to the excitation function used in physical models of musical instruments.

To simplify the configuration of molecules for audio synthesis, individual atoms can also be restrained to hold them in position using the spring force equation above with c set to a high value of 10000 kJ/mol and and $\vec{g}_i$ set to the initial position of the atoms.

The parameters of the molecular dynamics simulation can be exposed and controlled in real-time, allowing users to experience and explore a range of dynamical states. For example, the temperature T, friction γ and timestep t were exposed in this work (see Section 2.3.1). The simulation is configured using a Jupyter notebook in the Python language [21, 29], which is used to launch, configure and exit the simulations.

2.2 Molecular Scanned Synthesis

Molecular simulations are computationally very expensive and despite extensive GPU acceleration (with OpenMM), simulation data frames are processed at a rate of approximately 45 Hz[1]. Time-stepping the simulation at audio rates would incur an impractical computational cost, so we have used scanned synthesis to render audio rate streams from low rate simulation updates.

In Verplank's original scanned synthesis paper [41], the underlying physical model was a string, representing a one dimensional path that mapped directly onto a wavetable. This model was later extended by Boulanger et al [6] to higher dimensional structures with the recommendation that the "scanning trajectory should be moving between masses that are connected". In a molecular dynamics context, atoms can be considered analogous to masses and a scan path can be defined as a set of N contiguous atoms for a given molecule. The synthesis process then takes the form of wavetable synthesis where at each simulation time step, a wavetable is computed from the displacement of each scan path atom. Three Euclidean distance-based methods for computing atomic displacement have been developed:

(1) **Absolute displacement** involves computing the pairwise distance between the instantaneous position of each atom and its position at initialisation time.
(2) **Relative displacement** is taken as the distance between an atom's current position and its position at a previous time step. Previous atom positions are computed using recursive one-pole low-pass filters to ensure smooth convergence towards the current position.
(3) **Absolute centroid displacement** is intended for closed/circular scan paths beginning and ending at the same atom and involves calculating the displacement of each atom relative to the centroid of the scan path. This is achieved by storing the centroid of the scan path at initialisation time and then, at each timestep, subtracting the centroid distance of each atom from from its centroid distance at initialisation time.

The resulting wavetable represents an N element array which is scanned at the desired pitch (selected by MIDI) using liner interpolation. To avoid race conditions and discontinuities, a double buffering technique is employed whereby updated wavetable values are loaded into a temporary buffer on a background thread before being swapped into the audio thread for synthesis. To transition smoothly between successive wavetables, a second interpolation is used to fade between old and new wavetables, using the bilinear interpolation method described in [40]. To reduce aliasing, the wavetable is 8 × oversampled and then downsampled using 4-point Lagrange interpolation. A blocking filter is used to remove the DC component [35] before the signal is multiplied by an ADSR envelope generator to modulate the amplitude. The synthesis process is incorporated into a 32-note polyphonic sound engine written in C++ using the JUCE [34] framework to build a standalone application and a variety of plugin formats (including Unity and VST) to enable the synthesiser to run within a variety of environments and digital audio workstations (DAWs). The graphical user interface (GUI) is shown in Figure 2, incorporating a simple simulation visualiser, ADSR/simulation controls and a plot showing the current wavetable[2].

2.3 Control

Users can control the timbre of the molecular musical instruments in 3 ways: via the simulation parameters on the GUI, over MIDI/automation events, and by interacting with the simulated molecules using the VR client.

2.3.1 Simulation Parameters. The simulation parameters, temperature, friction and timestep, enable the dynamics of the simulation, and thus the timbre of the synthesiser, to be controlled from the

[1] running on a high-performance gaming laptop fitted with an NVIDIA GeForce RTX 2080 graphics card

[2] The project source code is available here: https://gitlab.com/teamaxe/GrapheneSynth

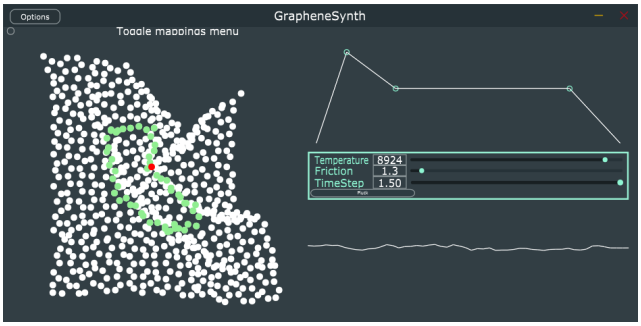

Figure 2: GUI for the graphene audio synthesis application, showing the atom positions with the scan path shown in green, initial centroid in red with a plot of the current displacement wavetable at the bottom

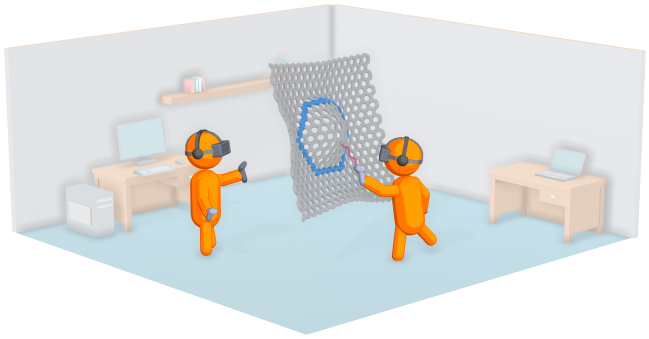

Figure 3: Example scenario with two users in VR manipulating the Graphene Sheet simulation. Credit: Alex Jamieson-Binnie, licensed under CC-BY-SA.

GUI. The temperature, in Kelvin, facilitates the heating or cooling of the simulation, causing atoms to vibrate faster or slower respectively. The effects of this parameter can be heard as an evolving texture which ranges from subtle to pronounced as the temperature is increased. Additionally, the atomic motion at higher temperatures results in energy perturbations permeating more rapidly throughout the structure of the simulated molecule. The friction parameter controls how aggressively the temperature is maintained with a high friction resulting in a highly damped system and, consequently, controls how rapidly the simulated molecule returns to its equilibrium state. Finally, the timestep parameter sets the speed of the simulation in femtoseconds, controlling how rapidly the timbre of the synthesiser evolves.

2.3.2 MIDI, Automation and Mapping. The simulation and ADSR parameters can also be controlled using MIDI control change message or by parameter automation within a hosting DAW. Pitch and gate are controlled by MIDI note on and off messages, which could originate from an external MIDI keyboard, sequencer or any other MIDI event emitter. Pitch is controlled by updating the internal wavetable interpolation function such that a single scan path/wavetable cycle matches the period of the selected MIDI note. Note events also trigger a molecular 'pluck', where a force is scaled in proportion to the note velocity and applied momentarily to the first atom in the scan path. The plugin also provides a GUI that allows users to map simulation features to MIDI and control change messages. This enables simulation features to modulate other plugins or any feature of a DAW that can be controlled using MIDI (see 4.3 for an example). Mapping from certain features, such as kinetic energy, presents a challenge as the value range is generally not known beforehand and will be unique to each different simulation. To ensure consistency, incoming values are normalised into a range computed by the mean +/− two standard deviations of the previous 20 readings.

2.3.3 VR. As an additional control option, the synthesiser timbre can also be modified via one or more connected Narupa VR clients, supporting multi-person interaction, see figure 3. As the visual representation of the molecular simulation is co-located with the interaction site in VR, users are able to apply forces to the simulated molecule directly. Pulling the trigger on a VR controller causes a spring force to be applied to the nearest atom j by setting $\vec{g}_i$ in equation 3 to the coordinate of the controller. This gives the effect of being able to grab and manipulate individual atoms (see Figure 5). Movements permeate though the molecular structure producing waves and ripples that that can be heard as fluctuations in the timbre of the synthesised sound.

3 SIMULATIONS

To explore the general applicability of the proposed method for sonifying/synthesising the structure of molecules, three molecular instruments were created using simulations of 17-alanine, graphene and a carbon nanotube. The geometry of each molecule requires a unique scan path comprising a contiguous sequence of atoms extending through the molecule's structure. The manually defined scan paths are described for each example below.

3.1 17-Alanine

17-alanine is a synthetic bio-molecule with a string like structure comprising a chain of 17 alanine amino acids. While this molecule is not naturally occurring, the alpha-helical structure of 17-alanine is a motif found in larger protein complexes, which are of particular significance to molecular scientists since they perform a vast array of functions that take place within living organisms.

The 17-alanine molecule (figure 4a) was simulated in Narupa with restraints placed on a pair of atoms at each end of the chain. The scan path is defined as a linear sequence along the *backbone* of the molecule, comprising a set of 50 adjacent atoms. The wavetable was computed using the absolute displacement method described in Section 2.2.

3.2 Graphene

Graphene is a form of carbon consisting of a single sheet of carbon atoms connected by covalent bonds in a repeating lattice of hexagons (see Figure 5). Graphene is intensely studied, both because it is a raw material for sculpting nanostructures such as carbon nanotubes and fullerenes and because, despite decades of theory stating that such 2D shapes would not be stable, it was found to be

stable and to exhibit desirable properties such as high strength, low density and efficient conductivity [14].

A 40 nm by 40 nm graphene molecule (less than 1000th the width of a human hair) was simulated. The graphene sheet was restrained at the corners and a hexagonal scan path comprising 54 neighbouring atoms was defined, as shown in figure 5. The wavetable was computed using the absolute centroid displacement method.

3.3 Carbon Nanotube

Carbon nanotubes also represent a hexagonal lattice of carbon atoms rolled up to form a hollow cylinder. Like graphene, these molecules are also widely studied as they are highly conductive both thermally and electrically, with exceptional tensile strength and profound technological implications.

A 60 atom nanotube was simulated (figure 4b) with a scan path defined as a spiral wrapping around the tube from one end to the other comprising all 60 atoms (see figure 4c). Since the unrestrained nanotube can freely drift around the simulation space the wavetable was computed using the relative displacement method.

4 EXAMPLES

This section presents four examples (three audio, one video) demonstrating the features of the system. Each audio example consists of a C3 note rendered with different parameter settings with spectrograms shown in figure 6. Examples were all created using Reaper [7] running the 17-alanine VST plugin and simulation. All media files are included with the supplementary materials[3].

4.1 Example 1: Static Parameters

In this example[4], a C3 note is played at full velocity, with no modulation. The note on event causes a single energy perturbation and the signal gradually fades away as the molecule returns to its equilibrium position.

4.2 Example 2: Automating Temperature

This example[5] uses the same values for friction and timestep as Example 1 but the temperature is modulated via the automation lane in Reaper. The sound begins with a temperature value of 0 K which rises to 5500 K between 0:02 and 0:09 and falls back to 0 K at 0:15. The higher temperature causes the molecule to rapidly vibrate and deform, corresponding to a brighter and noisier signal.

4.3 Example 3: Modulating External Audio Parameters

This example[6] uses a similar modulation trajectory to Example 2. However, it differs as the kinetic energy of the molecular simulation is emitted by the plugin as MIDI control change messages, which can be used to modulate audio parameters within host. In this case, kinetic energy is mapped to control the feedback gain and low-pass filter cutoff frequency of a native delay plugin within Reaper.

[3]https://zenodo.org/record/3965510
[4]https://soundcloud.com/user-688116459/mol-mus-instruments-example-1
[5]https://soundcloud.com/user-688116459/mol-mus-instruments-example-2
[6]https://soundcloud.com/user-688116459/mol-mus-instruments-example-3

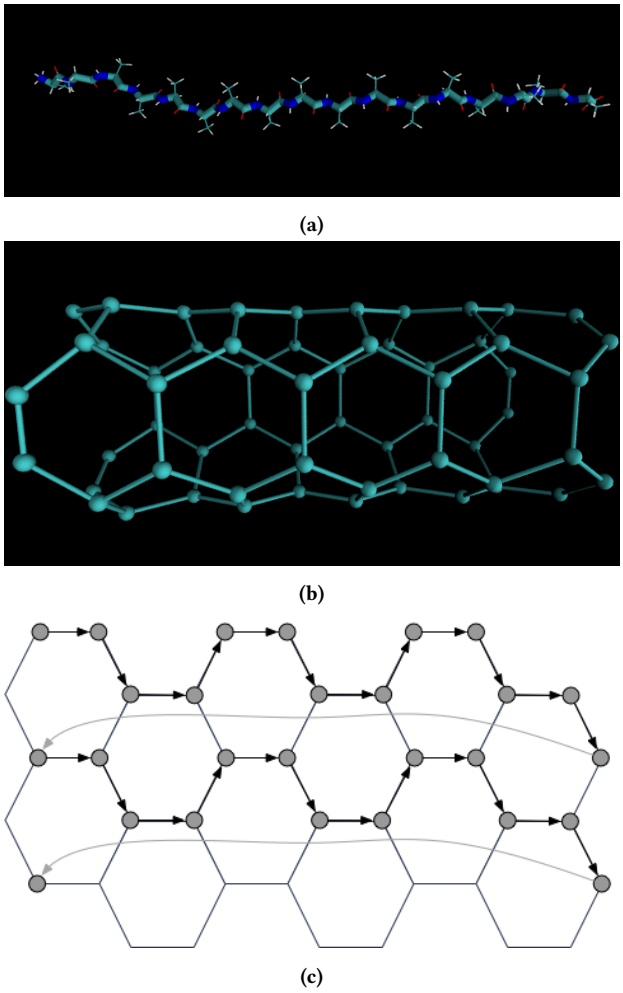

(a) **(b)** **(c)**

Figure 4: (a) shows the structure of 17-alanine. The scan path is shown as the thick tube running along its centre. (b) shows the carbon nanotube, which is formed of the same base hexagonal lattice as graphene. (c) depicts a 2D view of an unrolled section of the nanotube. Arrows indicate the direction of the scan path, resulting in a spiral traversing length of the tube.

4.4 Example 4: VR Interactions

This example video[7] shows both VR and audio clients connected to a 17-alanine simulation. Instead of using MIDI notes to trigger force perturbations, the VR client controls the movement of the molecule directly. The timbre of the pre-programmed chords can be heard to evolve as the VR operator interacts with the molecule. Faster movements generate louder tones with high frequency content and, conversely, the sound approaches silence when there are no interactions. The kinetic energy of the system is mapped to modulate the percussive sounds.

[7]https://vimeo.com/441255810

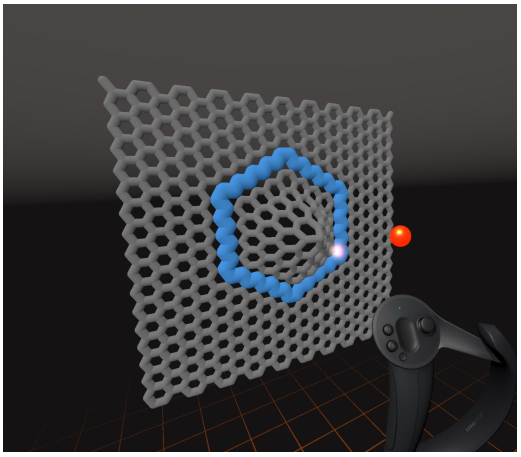

Figure 5: Graphene molecule simulated in the Narupa iMD application, with the trajectory used for synthesis highlighted and a user applying an interactive force with their VR controller.

5 CASE STUDY: COMPOSING WITH MOLECULAR MUSICAL INSTRUMENTS

In order to explore the aesthetic properties of the sonification method, and its potential utility in sound design and music, we used the graphene plugin to produce a short demonstration music composition[8]. The synthesis plugin (in real-time two way communication with the graphene simulation) was hosted within the Ableton Live DAW. Using the synthesiser as a plugin (rather than standalone application) brings advantages: multiple instances of the plugin can be used simultaneously, although with a single simulation running on the Narupa server. Plugin instances can also be placed within a network of audio routes and gain structures. Note events can be recorded live using a MIDI interface or using non-realtime sequencing techniques (both were used here). Similarly, automation events controlling plugin parameters can be entered using a mouse or MIDI interface. This automation can be used to control the molecular simulation at the same time as the simulation is controlling audio parameters. The synthesiser can of course be used in conjunction with other instruments and effects; however, to demonstrate the range of timbre produced by the synthesis method, all sounds in this piece originate from the graphene synthesiser alone with only limited use of Live's reverberation, delay and EQ. The aim of the composition was to explore the sonic range of the plugin. A brief analysis of the composition is provided here to highlight some of the techniques we evolved.

The piece opens with a slowly-evolving 'drone' section. The synthesis method lends itself to this sonic territory and consequently these sounds are relatively easy to produce. The timbral evolution of the technique was considered to sound most interesting at the high and low extremes of the MIDI range, and the introduction section to the piece tends to these extremes, with a low drone emerging before higher tones are introduced. Slow attacks and sustained envelopes are used, and gradual automation of the temperature, friction and

[8]https://soundcloud.com/user-688116459/mol-mus-instruments-graphene-music

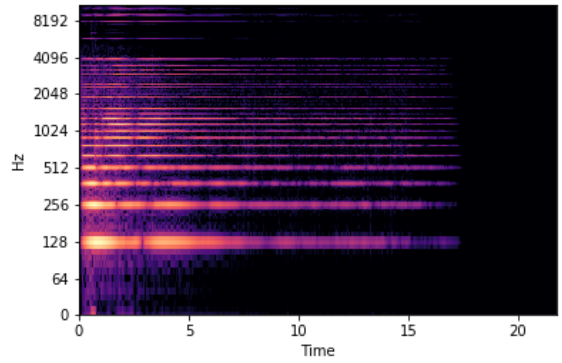

(a)

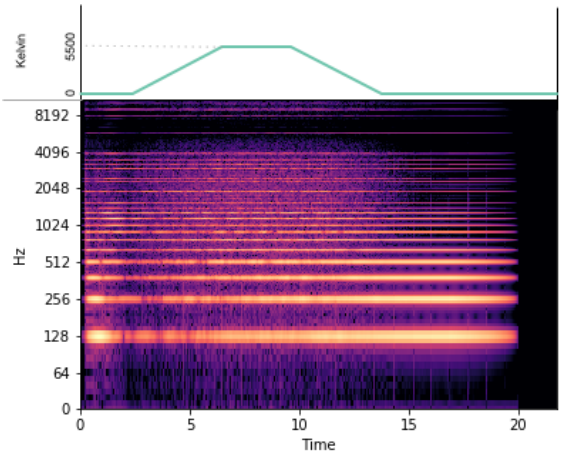

(b)

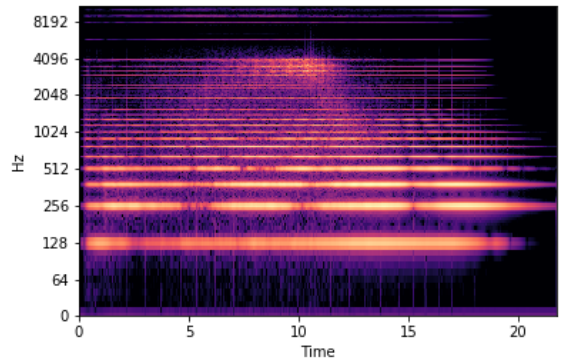

(c)

Figure 6: (a) Example 1 (Section 4.1) showing a single note spectrogram. (b) Example 2 (Section 4.2) showing a single note spectrogram below the temperature automation plot. (c) Example 3 (Section 4.3) showing a single note spectrogram with kinetic energy mapped to control a delay effect.

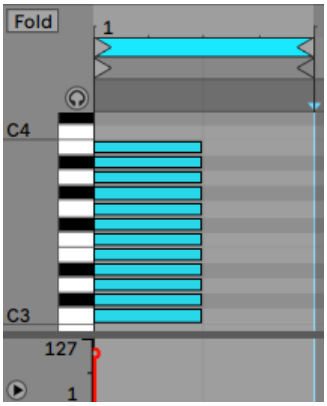

Figure 7: Image showing a 13 note cluster within the piano roll editor of Ableton Live

timestep of the simulation animates the sounds. The noisy 'windy' sound here uses a technique exploited throughout the composition to produce noisier and percussive sounds. A thirteen-note cluster including every semitone across an octave (see figure 7) was played with high temperature and timestep but low friction, producing something close to noise.

At ~0:30, mid-high metallic sounds are introduced. These also use the 13 note-cluster, but higher in the MIDI range and with a more pronounced envelope (short attack, long release tail) and a much higher friction value. This makes the pitch content more prominent and the sound more inharmonic than noisy. Throughout the whole of the first minute, the overall kinetic energy of the simulation is mapped to a low-pass filter on the main output. Since this energy is driven by the automated changes to temperature, friction and timestep, a feedback loop is produced between simulation and the sound, producing distinctive and slightly chaotic 'waves' in energy that are very evident both in the simulation and soundscape.

The rest of the composition is more controlled, structured and traditionally 'musical'. From ~1:00 we introduce percussive and textural sounds. These are again produced using a 13 note-cluster, transposed to various ranges, but this time with very short envelopes. Most of these sounds use almost the shortest attack and delay possible with sustain and release set to zero.

The final section, starting at ~2:00, introduces more rhythmically metric and harmonic material. The regular 'heartbeat'-like sound heard here is again produced using a 13-note cluster, but this time right at the bottom of the MIDI range. The 'sidechain compression' effect (which could have easily been made using that technique) is produced by automating the plugin parameters combined with a more controlled, faster version of the 'energy waves' used at the start.

6 CONCLUSIONS AND FUTURE WORK

In this paper we introduced a generalised approach for sonifying the dynamical data produced by a molecular dynamics simulation and explored its potential for audio synthesis. The details of the method are described which uses a scanned synthesis approach to translate low-rate simulation data into high-rate audio data. We use the method to sonify a range of simulated molecules including 17-alanine, graphene and a carbon nanotube. The effects of simulation parameters on the auditory character of the sound are examined and a case study and reflection on the use of the method within a music composition context is included. We are particularly interested to explore the musical application of the approach in order to understand the aesthetic factors that might enhance the usability and expressive quality of sonifications in molecular science [43].

The project currently builds to a variety of formats to enable the process to be easily used within both scientific and musical contexts. This includes the Unity Native Audio Plugin format, meaning that the work can be directly imported into the Narupa immersive molecular dynamics VR client, and used for scientific research [31]. The development and evaluation of the approach for molecular science will form the basis of future work.

In this paper the simulation host and audio client run on a single machine. However, the simulation, sonification and VR clients can also run on separate machines over local or wide area networks, enabling multiple participants to run, interact and listen to the simulation at the same time in different locations. The server can also run on the cloud, offloading the majority of the computational demands of the system and relaxing the requirement for high-specification GPUs.

The scan paths used for the audio synthesis process in this work have been manually defined. Exploring automatic methods for generating these paths will be an important next step. There are already a number of methods in molecular science for identifying features of proteins, such as primary/secondary structure or the backbone chain. We will explore a range of techniques for auto-generating scan paths in order to sonify arbitrary molecules.

ACKNOWLEDGMENTS

TJM, JH and DG acknowledge support from Leverhulme Trust, Royal Society, the British Academy and the Royal Academy of Engineering under Apex Award APX\R1\180118. TJM would also like to acknowledge support from the Bristol and Bath Creative Cluster (AHRC) and thank members of the CT Lab. AJJ is funded by PhD studentships from EPSRC. The authors would like to thank Joe Crossley-Lewis and Dr David Creasey for their input and advice.

REFERENCES

[1] Andrew Allen and Nikunj Raghuvanshi. 2015. Aerophones in Flatland: Interactive Wave Simulation of Wind Instruments. *ACM Trans. Graph.* 34, 4, Article Article 134 (July 2015), 11 pages. https://doi.org/10.1145/2767001

[2] Norman L Allinger, Young H Yuh, and Jenn Huei Lii. 1989. Molecular mechanics. The MM3 force field for hydrocarbons. 1. *Journal of the American Chemical Society* 111, 23 (1989), 8551–8566.

[3] Robert E. Arbon, Alex J. Jones, Thomas Mitchell, Lars A. Bratholm, and David R Glowacki. 2018. Sonifying Stochastic Walks on Biomolecular Energy Landscapes. In *Proceedings of the 24th International Conference on Auditory Display*. Houghton.

[4] Holger Ballweg, Agnieszka K Bronowska, and Paul Vickers. 2016. Interactive Sonification for Structural Biology and Structure-Based Drug Design. In *Proceedings of ISon 2016, 5th Interactive Sonification Workshop*. Bielefeld, 41–47. http://interactive-sonification.org/ISon2016/papers/ISon2016-06-BallwegBronowskaVickers.pdf

[5] Aldo Borgonovo and Goffredo Haus. 1986. Sound Synthesis by Means of Two-Variable Functions: Experimental Criteria and Results. *Computer Music Journal* 10, 3 (1986), 57–71. http://www.jstor.org/stable/3680260

[6] Richard Boulanger, Paris Smaragdis, and John Fitch. 2000. Scanned Synthesis: An Introduction and Demonstration of a New Synthesis and Signal Processing Technique. In *International Computer Music Conference*.

[7] Cockos. 2020. *Reaper Digital Audio Workstation.* https://www.reaper.fm (accessed: 03.07.2020).

[8] D. Creasey. 2016. *Audio Processes: Musical Analysis, Modification, Synthesis, and Control.* Routledge.

[9] Florian Dombois and Gerhard Eckel. 2011. Audification. In *The sonification handbook.*

[10] Peter Eastman, Jason Swails, John D. Chodera, Robert T. McGibbon, Yutong Zhao, Kyle A. Beauchamp, Lee-Ping Wang, Andrew C. Simmonett, Matthew P. Harrigan, Chaya D. Stern, Rafal P. Wiewiora, Bernard R. Brooks, and Vijay S. Pande. 2017. OpenMM 7: Rapid development of high performance algorithms for molecular dynamics. *PLOS Computational Biology* 13, 7 (07 2017), 1–17. https://doi.org/10.1371/journal.pcbi.1005659

[11] Sebastian L Franjou, Mario Milazzo, Chi-Hua Yu, and Markus J Buehler. 2019. Sounds interesting: can sonification help us design new proteins? *Expert Review of Proteomics* 16, 11-12 (dec 2019), 875–879. https://doi.org/10.1080/14789450.2019.1697236

[12] Daan Frenkel and Berend Smit. 2001. *Understanding molecular simulation: from algorithms to applications.* Vol. 1. Elsevier.

[13] Miguel Angel Garcia-Ruiz and Jorge Rafael Gutierrez-Pulido. 2006. An overview of auditory display to assist comprehension of molecular information. *Interacting with Computers* 18, 4 (2006), 853 – 868. https://doi.org/10.1016/j.intcom.2005.12.001 Special Theme Papers from Special Editorial Board Members (contains Regular Papers).

[14] Andre K Geim and Konstantin S Novoselov. 2010. The rise of graphene. In *Nanoscience and technology: a collection of reviews from nature journals.* World Scientific, 11–19.

[15] Google. 2020. *GRPC.* https://grpc.io (accessed: 03.07.2020).

[16] Florian Grond and Jonathan Berger. 2011. Parameter mapping sonification. In *The sonification handbook.*

[17] Florian Grond and Fabio Dall'Antonia. 2008. Sumo. A Sonification Utility for Molecules. In *Proceedings of ICAD 2008.* International Community for Auditory Display, 1–7. https://smartech.gatech.edu/handle/1853/49952

[18] Florian Grond and Thomas Hermann. 2014. Interactive Sonification for Data Exploration: How listening modes and display purposes define design guidelines. *Organised Sound* 19, 1 (2014), 41–51.

[19] Florian Grond, Stefan Janssen, Stefanie Schirmer, and Thomas Hermann. 2010. Browsing RNA Structures by Interactive Sonification. In *Proceedings of ISon 2010, 3rd Interactive Sonification Workshop.* Stockholm, 11–16. http://www.interactive-sonification.org/ISon2010/proceedings/papers/GrondJanssenSchirmerHermann{_}ISon2010.pdf

[20] Roald Hoffmann. 1990. Molecular Beauty. *The Journal of Aesthetics and Art Criticism* 48, 3 (jan 1990), 191–204. https://doi.org/10.2307/431761

[21] Thomas Kluyver, Benjamin Ragan-Kelley, Fernando Pérez, Brian E Granger, Matthias Bussonnier, Jonathan Frederic, Kyle Kelley, Jessica B Hamrick, Jason Grout, Sylvain Corlay, et al. 2016. Jupyter Notebooks-a publishing format for reproducible computational workflows.. In *ELPUB.* 87–90.

[22] Gregory Kramer. 1994. Some organizing principles for representing data with sound. *Auditory Display-Sonification, Audification, and Auditory Interfaces* (1994), 185–221.

[23] Intangible Realities Laboratory. 2020. *Narupa iMD Code Repository.* https://gitlab.com/intangiblerealities/narupa-applications/narupa-imd (accessed: 03.07.2020).

[24] Intangible Realities Laboratory. 2020. *Narupa Protocol Code Repository.* https://gitlab.com/intangiblerealities/narupa-protocol (accessed: 03.07.2020).

[25] Ask Hjorth Larsen, Jens Jørgen Mortensen, Jakob Blomqvist, Ivano E Castelli, Rune Christensen, Marcin Dułak, Jesper Friis, Michael N Groves, Bjørk Hammer, Cory Hargus, Eric D Hermes, Paul C Jennings, Peter Bjerre Jensen, James Kermode, John R Kitchin, Esben Leonhard Kolsbjerg, Joseph Kubal, Kristen Kaasbjerg, Steen Lysgaard, Jón Bergmann Maronsson, Tristan Maxson, Thomas Olsen, Lars Pastewka, Andrew Peterson, Carsten Rostgaard, Jakob Schiøtz, Ole Schütt, Mikkel Strange, Kristian S Thygesen, Tejs Vegge, Lasse Vilhelmsen, Michael Walter, Zhenhua Zeng, and Karsten W Jacobsen. 2017. The atomic simulation environment—a Python library for working with atoms. *Journal of Physics: Condensed Matter* 29, 27 (2017), 273002. http://stacks.iop.org/0953-8984/29/i=27/a=273002

[26] Ryan McGee. 2013. VOSIS: a Multi-touch Image Sonification Interface.. In *NIME.* 460–463.

[27] Ryan McGee, Daniel Ashbrook, and Sean White. 2012. SenSynth: a Mobile Application for Dynamic Sensor to Sound Mapping.. In *NIME 2012.*

[28] Thomas Mitchell, Joseph Hyde, Philip Tew, and David R Glowacki. 2014. dance-room Spectroscopy: At the Frontiers of Physics, Performance, Interactive Art and Technology. *Leonardo* 49, 2 (aug 2014), 138–147. https://doi.org/10.1162/LEON_a_00924

[29] Thomas J. Mitchell and Alex J. Jones. 2020. *Graphene Synth Repository.* https://bitbucket.org/teamaxe/graphenesynth (accessed: 03.07.2020).

[30] Donald A Norman. 2004. *Emotional design: Why we love (or hate) everyday things.* Basic Civitas Books.

[31] Michael B. O'Connor, Simon J. Bennie, Helen M. Deeks, Alexander Jamieson-Binnie, Alex J. Jones, Robin J. Shannon, Rebecca Walters, Thomas J. Mitchell, Adrian J. Mulholland, and David R. Glowacki. 2019. Interactive molecular dynamics in virtual reality from quantum chemistry to drug binding: An open-source multi-person framework. https://doi.org/10.1063/1.5092590 arXiv:1902.01827

[32] Joshua Peschke and Axel Berndt. 2017. The geometric oscillator: Sound synthesis with cyclic shapes. In *ACM International Conference Proceeding Series*, Vol. Part F1319. Association for Computing Machinery. https://doi.org/10.1145/3123514.3123522

[33] Benjamin Rau, Florian Frieb, Michael Krone, Christoph Muller, and Thomas Ertl. 2016. Enhancing visualization of molecular simulations using sonification. In *2016 IEEE 1st International Workshop on Virtual and Augmented Reality for Molecular Science (VARMS@IEEEVR).* IEEE, Arles, France, 25–30. https://doi.org/10.1109/VARMS.2015.7151725

[34] Roli. 2020. *The JUCE Library.* https://juce.com (accessed: 03.07.2020).

[35] Julius.O. Smith. 2010. *Physical Audio Signal Processing, online book, 2010 edition.* W3K Publishing.

[36] Robert D Sorkin. 1988. Why are people turning off our alarms? *The Journal of the Acoustical Society of America* 84, 3 (1988), 1107–1108.

[37] Marc Sosnick and William Hsu. 2011. Implementing a Finite Difference-Based Real-time Sound Synthesizer using GPUs. In *Proceedings of the International Conference on New Interfaces for Musical Expression.* 264–267. https://doi.org/10.5281/zenodo.1178173

[38] Mark D Temple. 2017. An auditory display tool for DNA sequence analysis. *BMC bioinformatics* 18, 1 (2017), 1–11.

[39] Noam Tractinsky. 2004. Toward the study of aesthetics in information technology. *ICIS 2004 proceedings* (2004), 62.

[40] Robert Tubb, Anssi Klapuri, and Simon Dixon. 2012. The Wablet: Scanned Synthesis on a Multi-Touch Interface. In *International Conference on Digital Audio Effects.*

[41] Bill Verplank, Max Mathews, and Rob Shaw. 2001. Scanned Synthesis. *The Journal of the Acoustical Society of America* 109, 5 (2001). http://www.billverplank.com/ScannedSynthesis.PDF

[42] Paul Vickers. 2006. Lemma 4: Haptic input+ auditory display= musical instrument?. In *International Workshop on Haptic and Audio Interaction Design.* Springer.

[43] Paul Vickers. 2016. *Sonification and Music, Music and Sonification.* Taylor Francis, United Kingdom, 135–144.

[44] Chi-Hua Yu, Zhao Qin, Francisco J. Martin-Martinez, and Markus J. Buehler. 2019. A Self-Consistent Sonification Method to Translate Amino Acid Sequences into Musical Compositions and Application in Protein Design Using Artificial Intelligence. *ACS Nano* 13, 7 (jul 2019), 7471–7482. https://doi.org/10.1021/acsnano.9b02180