1 The Sonification Continuum

Music is continuous, it is only listening that is intermittent.

John Cage[1]

1.1 Prelude: Background

The craft of collecting objective, measurable facts and information ('data') and constructing systematic processes to represent them in sound ('sonification') has been a central part of my compositional practice for over twenty years. This approach has produced hundreds of audio and multimedia outputs, collaborative projects, presentations, masterclasses, broadcasts and events with a wide range of scientists and institutions. Despite this level of engagement, it has never been an approach I was introduced to formally or guided through. However, when looking back it appears to have evolved along an unplanned path guided by landmarks and faint signposts. Some of these were moments of diverting curiosity and others almost overwhelming personal experiences that somehow called me to engage with sonification, but with rigour, honesty and sincerity. These experiences were essentially my teachers in the craft and include the following:

(a) An early life at CERN (Geneva, Switzerland), where scientific curiosity and a fresh and total musical immersion were unified by a shared sense of discovery and wonder.

(b) An exhibit at CERN, where passing photons which happened to drift and end their light-years of travel in the vicinity emitted a click in a satisfying and mesmeric ungridded rhythm.

(c) Childhood summers in Greece, staring at the night sky and transcribing a melodic contour in the stars.

(d) The discovery of computer music, the MIDI and DAW environments, and generative music where a nexus of simple objective instructions led to a richness of subjective musical experience.

(e) Encountering the practice of musical cryptogram in the BACH and Shostakovich motifs, alongside the *milimetrazação* (graphing or millimetrisation)[2] technique of Villa-Lobos's *New York Skyline Melody* – whose melody (and architecture) I know as well as any conventionally conceived melodic contour.[3]

---

[1] John Cage cited in Millar and Cage, 2010:74.

[2] See Enyart,1984:188 and Slonimsky, 1945:6.

[3] For a visualization we produced of Villa-Lobos's *New York Skyline Melody,* see Tanczos (2014).

(f) First hearing the sonified output of my engagement with Rilke's thought experiment (see Section 3.1.1 Primal Sound) – a piece of music (which has had a life beyond its origin) that I neither had expected from the process nor could otherwise imagine.

(g) Sonifying my daily blood results from my hospital bed while undergoing treatment for leukaemia (see Section 3.1.2 Bloodlines).

(h) Experiencing a reaction to a medication which caused the pitch Bb2 to create an intense frisson experience. The effect was so pronounced and repeatable that, despite not having perfect pitch, I could – using the 'chill level' as an analytical tool – transcribe chords and melodies.

(i) Using our yet-to-be daughter's embryo scan to generate a piece of music[4] – an ode using objectively sourced pitch and rhythmic data despite the uncertainty of her future existence.

(j) The cascade of data-musical insights, revelations from the Sound Asleep project (Section 3.1.5).

(k) Deciding on the morning of a TEDx presentation to translate the audience selecting their seating positions into sound, conceiving and enacting the system and presenting it to them at the end of my talk (Section 3.1.8).

(l) Running a digital image of Monet's *Water Lilies* through my newly created Kandinsky patch – a piece of coding that translates colour data into discrete and continuous pitch material (see Section 3.1.9).

Experiences such as these are rich and integrate my musical and personal life. And somehow, when I belatedly arrived at the academic literature – as thoughtful and useful as it is – it did not resonate fully with my experiences and fundamental motivations. The perspectives I encountered tended to present either somewhat prescriptive and proscriptive definitions of data sonification, or 'anything goes' compositional tools where data is selected, reworked, manipulated and even abandoned in service of a musical end. In the domain of 'pure' data sonification, music is treated with suspicion as a potential distortion of data communication, where composers cannot be trusted to resist aesthetic urges rather than staying 'true to the data'. This wariness towards musical intent in the context of data sonification is apparent even in key academic music texts: *The Oxford Handbook of Computer Music* has a chapter on data sonification but 'data music' (Worrall, 2011) is addressed only in a final subsection, where it exists on a 'continuum' of sonification engagement. More tellingly, the chapter on data sonification in *The Oxford Handbook of Algorithmic Music* is entitled 'Sonification ≠ Music' (Scaletti, 2021), laying bare the supposed distinction. Such positions contain valid

---

[4] *Two Blue Circles for Classical Guitar and Electronics* (Mermikides, 2020b).

underlying concerns, particularly Neuhoff's (2019) clear identification of the challenges and potential limits of sonification. I too favour clear, unfudged data – and scientific – communication, and I strongly advocate prioritising data communication over 'sonic entertainment' to overcome barriers of literacy, numeracy and visualisation (see Sawe et al., 2020) and to enhance physical development and rehabilitation (see for example, Scholz et al., 2016). And yet to separate data communication from musical mechanics appears to me to make an unhelpful opposition, for as I shall discuss later, it rests on assumptions on the definition and limits of 'music'. I have since found some kinship in the sonification works of John Luther Adams, Xenakis, Cage and members of the International Community for Auditory Display (ICAD), but equally, I have recognised traces of sonification in many 'conventional' works, and these traces have proved just as stimulating.

There are others who allow scientific and artistic sonification to have hybrid goals, as in Gresham-Lancaster's warning that 'science must be inclusive of craftsmanship and artistry, and vice versa, for this field to be fully accepted and realise the usefulness and promise of these important tools' (Gresham-Lancaster, 2012:212). Nonetheless, the distinction between data sonification and a loose form of data music (where the data is often seen as inspirational 'source' material rather than central to the music itself) still persists as in Neuhoff's 'bifurcation' of 'empirical' and 'artistic sonification' and his characterisation of any hybrid goals as a 'muddled middle' (Neuhoff, 2019). This wariness of the middle ground between the artistic and empirical aims of sonification may be well motivated, but it rests on narrow definitions of both music and communication of information, compromising two immediate scientific and artistic opportunities. Firstly, data communication can and should exploit the wealth of compositional tools and strategies available. Secondly, the challenge of representing data fairly tests and extends a composer's skills and potential outputs; it exposes hidden preconceptions about music while revealing its deep possibilities. If we accept music itself as a form of sonic communication – that is, the expression in sound of such information as patterns, processes, thoughts, narrative structures and states of emotion, mind or place, where sound includes speech, natural sounds, sound design, noise and pitch – then it becomes less clear, less useful to draw a hard line between data sonification and music.

Thus, I have never thought of data sonification (and/or data music)[5] as the novel and fleeting diversion from conventional music-making that it appears to

---

[5] Definitions and delineations of data sonification, data music, sonification and other terms are addressed directly in Section 1.2.

be with others, but as a fundamental part of my musical purpose. My approach often contests the boundaries between data sonification and broadly defined composition. When I have felt a composition to be successful, it has been because the data and its inherent patterns did more than provide a creative constraint or tribute; rather, they *mattered* in terms of sonic output, genuinely dictating and informing fundamental aspects of the compositional process and the experience of the listener where the data and its inherent patterns not only provide a creative constraint and tribute to the work, but *matter* in terms of sonic output. In collating and exhibiting some of my works in 2007, I hurriedly coined the moniker *Hidden Music*. This phrase, which implies that there is a music *inherent* in the data waiting to be unearthed, rather than music made from reshaping 'data material', has proved reliably and increasingly apt, and so I have retained it as a title for this Element. In these pages, I aim to illuminate the intersection of sonification and compositional practice, and to complement the existing literature for practitioners, science communicators and those interested in this young and promising field. In fact, the deeper my exploration of this craft, the more convinced I become that there is a complementary relationship between conventional musical theory – *the nature of music* – and data sonification – *the music of nature*. The former involves identifying and manipulating objective data in music, while the latter involves the sonic representation of objective data. In either domain, analytical models, concepts and technologies may be routinely flipped in order to pass material from one to the other – from data to music or music to data.

Despite the academic context of this Element, it aims to explore – rather than dictate – relevant ideas and approaches in a manner accessible and practical to composers, theorists and others. A history and/or survey of contemporary sonification practice is beyond the scope of this Element, particularly given problems of definition and boundaries.[6] Instead, I present conceptual and practical ideas to help the reader chart their own course, with selections of my work serving as examples (and not exemplars) of the underlying themes. Section 1 is concerned with first principles, addressing the complex of definitions, challenges and opportunities – essentially, the *what* of sonification. Section 2 confronts the *how* of sonification: principles, strategies, techniques and technologies of translations from the data to the sonic realm. Section 3 presents a selection of my works from the past two decades in order to illustrate these approaches in action, followed by some thoughts on the – or at least my – *why* of sonification.

---

[6] For an excellent introduction and overview of empirical sonification practice, see Supper (2016).

1.2 Unweaving Definitions

This Element is concerned with the set of practices that involve the communication of data through sound (including music). Depending on one's definition of data and communication, questions of boundary arise with the representation of any extra-musical content in music – be it the sophisticated language of talking drums in sub-Saharan Africa (Gleick, 2011:18–23), the sonic lightning bolts in Beethoven's Sixth Symphony, Joni Mitchell's ad hoc guitar tunings to 'the crows and the seagulls, and sonic references available' (Mitchell, 1994: 0:32–38), Villa-Lobos's use of the *Belo Horizonte* mountain skyline as thematic material in Symphony no. 6 (Enyart, 1984:188), Aphex Twin's embedding of images in the spectrographs of electronic works (see Buckle, 2022), the Tuvan *borbannadir* throat singing technique of closely imitating bubbling brooks (Aksenov, 1973) or any other innumerable examples. Even with more constrained definitions, there are numerous terms for the craft including sonification, data sonification (*'datason'*), audification, musical translation, auditory display, musification, data music, and data-, data-driven or data-based composition. This proliferation of variously defined and overlapping terms emerges not just from the rapid and unfettered evolution of this type of practice in an increasingly 'datarised' world, but from questions of intents, motivations, adherence to various systems of data selection and collection, and the ultimate use and reception. Some commenters have attempted to strictly define and demarcate these terms, most particularly in the delineation of a strictly defined data sonification (or simply 'sonification') from other artistic practices (see, for example, Hermann, 2010; Barrass & Vickers, 2011; Neuhoff, 2019; and Scaletti, 2021).

A commonly referenced description in the literature is 'the use of non-speech audio to convey information' (Kramer et al., 1999:1). Other definitions share a similar basic structure: a communication of information in sound, but with a caveat, the exception of the 'normal' sonic communication of speech, a 'seeing with our ears' (Vickers, 2016: 135). While some definitions focus on a systematic and reproducible *process*, others add a condition of the *purpose* of the process, such as a "mapping of numerically represented relations in some realm under study" to relations in an acoustic realm *for the purpose of interpreting, understanding, or communicating relations in the domain under study*" (italics added) (quoted in Barrass & Vickers, 2011:147) or Worrall's 'acoustic representation of data for relational interpretation by listeners, *for the purpose* of increasing their knowledge of the source from which the data was acquired' (Worrall, 2009: 314 – italics added). So while some sort of translation of data into sound is a basic requirement, the type (and method of collection) of

data, translation process, level of intervention or faithfulness and the purpose of the sonification, provide for various commenters additional criteria to this core requirement.

Though my work in this field happens to align with a more faithful and systematic representation of data than many of my sonification-curious colleagues, I find Baxter's perspective of treating sonification 'more of a verb rather than a noun, a technique to be used, rather than a thing to be arrived at' (Baxter, 2020:16) to be a helpful parry to the challenges of precise definition. Nonetheless, it is illuminating and helpful to interrogate our definitions of (data) sonification, and any hidden presumptions, inconsistencies and implications involved, which we now explore.

1.3 The Boundaries of Sonification and the Music of Music

Rather than rely dutifully on pre-existing definitions, I invite the reader to develop their own understanding and categorisations of the field. As a necessary (but not perhaps sufficient) condition for discussion, let us agree that all sonification entails a translation of some information into sound. The encoding of information and its transmission to a recipient invites engagement with Shannon's information theory, entropy and cryptography.[7] How the information is transmitted, the 'key' to its encoding, the degree to which its entropy (or conversely, its level of order) is preserved amidst 'noise', and whether the received signal can be decoded to reveal the original message, is a simple but clear framing of the challenge and craft of data sonification. Incidentally, this description applies to music in general, if one accepts that information includes the meaningful brain activity of emotional states (see Stark, Vuust & Kringelbach, 2018). If we understand *meaningful* in its usual senses of valuable, recognisable and not directly expressible, and if we take the *sender* and *receiver* to refer to human brains and *sound* as the transmission channel through which the signal is sent, then we arrive at a description of musical communication I am happy to accept.

So far, then, we have said that data sonification involves a transfer of information – meaningful signals – from sender to receiver, via some form of encoding. Unlike in conventional music, these meaningful signals do not necessarily originate from musical ideas or a complex of patterns in human brains, but from external data. We might picture the sonification process as a membrane through which data (observed properties in the domain under study) pass through to become sound, or 'readily sounded' material such as captured audio or musical instructions. This membrane divides the *data realm*

---

[7]   See, for example, Shannon, 1948: 379–423, 623–656 and Gleick 2011: 191–217.

(observed information) from the *sonic realm* (sounding or readily sounded objects), even as it provides access to both.

What constitutes material in the data realm? An essential definition may be 'observed facts and information'; however, the process of collection (scientifically objective or otherwise) and intended use (e.g. information collected for reference or analysis, etc.) of the data may differ between individuals, as in Scaletti's 'purposes' earlier. What does seem to be fundamental for data sonification (and sonification more broadly) is that there is a significant degree of objectivity in the data set; the data is collected (or to some degree selected) rather than created specifically for the purpose of sonification. If the data is mined, manipulated or ignored freely, then the process becomes closer to conventional composition, albeit through an elaborate pantomime: the system is played by a human like an instrument, rather than having the sonic output dictated by the observed material. All may be fair in composition, some ends may justify some means, and the piece may not exist at all without the original impetus. And yet such a process, I argue, cannot be called sonification. Even with artistic or loose approaches, sonification entails a delegation of decisions to the data: no matter how intricate the hands-on system design, there is a significant hands-off moment, where the system is allowed to run free of intervention. How this delegation might be assessed, and the extent to which it reaches stricter definitions of data sonification, are explored further in Sections 1.4 and 2.1.

So the manner of data collection might affect one's assessment of what type of sonification a process is – or whether it is sonification at all. But what of the data itself? Is any type of data – if collected objectively – allowable? Suppose we set up a technological device to somehow sample rapidly (over 40,000 times a second, say) changing air pressure, convert and store this information to a series of discrete values and then later reverse the process, using the captured data to recreate similar changes in air pressure. Is this sonification? This is, of course, a mischievous description of digital audio recording which – despite being a very clear example of objectively collected data being sounded – is excluded from conventional definitions of data sonification. Is a flute performance a sonification of the flautist's manipulation of air pressure and finger movements – to say nothing of the musical symbols on the page or the composer's abstracted ideas? There is no end to such Socratically awkward examples: imagine placing a sensor under each key of a piano keyboard which measures when, how long and how fast (or hard) each key is pressed, and this data is stored numerically and then used to replay digital audio samples of a piano at analogous pitch frequencies, durations and velocities. Is this a sonification of the keyboardist's performance, even if we systematically

manipulate the data (via transposition, scaling the velocity, doubling at octaves or triggering alternate instruments)? These descriptions of digital audio, a flautist's performance and a MIDI instrument are all clear examples of 'objective data' being 'mapped to sound' in a way that can 'convey information' of the 'domain under study' (all descriptions found in Section 1.2). Why then would these *not* be defined as sonification (data or otherwise) or data music? And why, for that matter, is speech – perhaps the clearest example of information being imparted through sound – explicitly excluded from many definitions of sonification? One possible response has been suggested, namely, that when data is generated to produce a particular sound, the resulting sound cannot be considered a sonification. We should be careful with this line of reasoning, however, for we will later see examples of how musical data can be sonified. Another possible response is that they are indeed examples of sonifications, but so established, conventional or mundane that they seem fundamentally different from the more novel auditory displays.[8] But our model of data and sonic realms offers an explanation that I prefer: that digital audio, musical notation, MIDI information and so on – although they are expressible as data – can be thought to already lie in the sonic realm in our model (which includes sound or material that is readily sounded). Turning this 'sonic data' into sound is, I suggest, usually a manipulation (recording, performance, editing, sequencing, transcription) *within* the sonic realm rather than a cross-realm translation through the sonification boundary. What seemed a simple description of data sonification (data translated into sound) becomes a rather puzzling recursive cycle of questions. We might also suggest why, for example, speech – and implicitly audio, notational and other 'sonic' forms of information – is explicitly excluded from common descriptions of sonification: speech is an established sonification practice, while the others already inhabit the sonic realm.

Figure 1 illustrates a number of these relationships. A sonification membrane separates the data and sonic realms. But this boundary is porous: material can travel from the data to the sonic realm using various techniques(see Section 2.3).

The sonic realm might be further subdivided into four domains:

1. The *audio domain*, which includes stored or transitory analogue and digital audio material.
2. The *instrumental domain*, which includes acoustic and electronic instruments, sequencers and their human and machine performers.

---

[8] Scaletti (2021) attests that music is in fact a sonification – of musical thought, a concept that is discussed later. Readers unfamiliar with this field's terms of art might be struck by the novel use of *display* to refer to something heard rather than seen; but to display – etymologically, 'to unfurl' – data in sound is surely an apt metaphor for the sonification process.