Rethinking the BEAST: Recent developments in multichannel composition at Birmingham ElectroAcoustic Sound Theatre

SCOTT WILSON* and JONTY HARRISON**

Birmingham ElectroAcoustic Sound Theatre, Music Department, University of Birmingham, Edgbaston, Birmingham B15 2TT, UK
E-mail: *s.d.wilson.1@bham.ac.uk; **d.j.t.harrison@bham.ac.uk

This paper outlines some recent developments in multichannel composition at the Electroacoustic Music Studios, University of Birmingham and its performance wing, BEAST (Birmingham ElectroAcoustic Sound Theatre). In doing so it attempts to codify and define some emerging aspects of spatialisation practice which are found both within the BEAST community and beyond. The discussion covers software and techniques developed and adapted for use with BEAST, new and pragmatic approaches to composing for large-scale multichannel systems, such as n-channel composition and composing in 'stems', and issues arising from a resulting blurring between composition and performance practices.

1. INTRODUCTION

The landscape of multichannel concert presentation of electroacoustic music has changed radically in recent years. Of primary importance in this change is the new availability of relatively inexpensive commercially produced multichannel hardware. This has considerably increased the possibilities and implications of working in multichannel formats, regardless of the style or design of one's chosen presentation system.

This process began perhaps in the 1990s with the introduction of digital multichannel recorders – such as the Alesis ADAT. These devices allowed for the 'eight-channel ring piece' to flourish as a semi-standard format (along with at least two common loudspeaker layouts; referred to within BEAST colloquially as the French and American configurations: four left-right pairs vs. a quadrophonic + interspersed offset quad (or 'double-diamond') respectively),[1] allowing for an easy exchange of works between diverse research institutions, artists and concert-presenting organisations. This has arguably led to refinements in both the technical and aesthetic aspects of multichannel audio art, and – along with other related developments, such as the increased availability of multichannel home theatre systems and powerful laptop computers – has had a democratising effect on the field as a whole.

In more recent years this development has accelerated through the availability of multichannel computer audio interfaces such as the Mark of the Unicorn 24I/O.[2] Such technology has made flexible digital software-based presentation systems possible without requiring custom designed and built hardware, and thus allowed for a move away from less flexible analogue setups which form the basis of traditional stereo diffusion systems. These developments bring with them considerable new prospects in multichannel composition and system design. These nascent possibilities require new strategies and aesthetic considerations, and have implications for presentation, performance and reception. How best can one make use of the flexibility of these large-scale digital systems? Since these systems make large-scale multichannel work – defined here for convenience as anything greater than eight channels – more readily possible, what technical and compositional issues need to be addressed, and which solutions will be most successful? More practically, in what fashion should composers, construct, store and transmit their works?

Whilst addressing these questions comprehensibly would a daunting task, this article will nevertheless attempt to explore some possible answers, and look at the solutions and approaches tried with the BEAST sound system at the University of Birmingham, hopefully offering some useful advice based on experience gained 'on the ground', through putting on concerts with a working large-scale multichannel loudspeaker system.

[1]As composer Eric Lyon has said, 'Eight channel is the new stereo'. As should be obvious from the two common configurations noted above, however, transportability cannot be guaranteed, even within a supposedly simple, straightforward definition of 'eight-channel' – as anyone attempting to play a work composed for the 'French' eight-channel array on an 'American' double-diamond system (or vice versa) can testify! Other idiosyncratic eight-channels configurations exist, and thus while eight channels as a media format may have become relatively standard, as a speaker configuration it does present some compatibility issues.

[2]As a tangible example, a single computer can straightforwardly control 4 24I/Os as a single audio interface, allowing 96 simultaneous inputs and outputs. Larger configurations are possible.

Organised Sound **15**(3): 239–250 © Cambridge University Press, 2010. doi:10.1017/S1355771810000312

https://doi.org/10.1017/S1355771810000312 Published online by Cambridge University Press

2. ABOUT BEAST

Founded in 1982 by Jonty Harrison, Birmingham ElectroAcoustic Sound Theatre is a large, non-homogenous loudspeaker presentation system. In its current incarnation we are capable of mounting systems of sizes in excess of 100 loudspeakers (circa 80 is typical), each addressable as a discrete channel. While we would refer readers elsewhere for a more detailed discussion of BEAST's design principles and historical development (Harrison 1988, 1998, 2000), it can loosely be described as a system which developed on the French acousmonium model, its primary purpose being live diffusion of stereo acousmatic works.[3] This remains a central aspect of BEAST activity, albeit one that increasingly coexists with other approaches and configurations – in other words works requiring different basic speaker configurations such as eight-channel rings (using both the 'French' and 'American' orientations), 5.1, n-channel (see below), ad hoc, and so forth – and at times supporting other musical idioms and genres, such as mixed instrumental and electroacoustic works, live electroacoustic performance, and multimedia.

While it is not our intention here to take sides in the technical and aesthetic battles over approaches to electroacoustic presentation (a debate in which – given the necessarily hybridised nature of the BEAST system – we are not inclined to be particularly partisan in any case), as a prelude to discussing more recent developments at BEAST we would like to dwell briefly upon what we feel are two of the more salient virtues of past BEAST practice, and of stereo diffusion as a presentation strategy in particular: adaptability and pragmatism.

3. ADAPTABILITY

As a strategy stereo diffusion is particularly flexible, both in terms of allowing for works to be adapted to available systems and for systems to be designed to make the most of available spaces. Since a piece intended for stereo diffusion is not in principle tied to a particular system configuration or localisation scheme (although the diffuser and/or composer can certainly have a general or specific one in mind) it can easily be adapted to make the most of what is available in a given situation.

Similarly, systems can be designed to take advantage of the idiosyncrasies and variations found in different performance spaces, even to the extent of making problems or limitations into opportunities waiting to be exploited. As noted above, BEAST is a 'non-homogenous' system – meaning that its various pairs and arrays of speakers are not all of the same size, type, or model – and this is a classic example of making a virtue of what for some might be necessity. Such a configuration allows for things like variations in character, and naturally facilitates 'spectral splitting' effects (see below), which can increase an audience's sense of envelopment by enhancing the diffuseness of sound materials.

Large-scale multichannel composition raises its own issues with adaptability (providing one wishes one's works to have any portability at all), given that there is little if any standardisation in terms of layout, hardware or software amongst most such systems currently in use. Adaptability is of course both an advantage and a requirement, and adapting a system to unfamiliar spaces necessitates adapting pieces to an (at least slightly) unfamiliar system. This requirement is by no means a new one from our perspective, however. BEAST has throughout its history been a touring system, relatively if not completely unique amongst large-scale loudspeaker systems in this regard, and maintaining a completely consistent setup from venue to venue would not only be 'missing a trick' in terms of the possibilities each space presents, it would be quite simply impossible in many cases. Thus even in terms of 'traditional' stereo diffusion alone it has been necessary to conceive of works as adaptable to different presentation systems to a significant extent. The lack of any standardisation amongst the variety of large- and small-scale multichannel systems in the world today and the types of strategies required for dealing with this situation are simply a continuation of this trend, albeit one which at times entails resolving somewhat more difficult technical complications.

4. PRAGMATISM

Pragmatism goes hand in hand with adaptability, and embracing the latter at every stage of the compositional process from conception to presentation is certainly an example of the former.

In the discussion of the non-homogenous nature of the BEAST system above, the notion that loudspeakers cannot be treated as strictly neutral and transparent conveyors of fully and ideally realised sound material is implicit. From a pragmatic rather than idealistic point of view this lack of neutrality is not a weakness but is again rather a potential tool waiting for an appropriate opportunity to be exploited.

One can see other pragmatic aspects in the typical BEAST arsenal of diffusion techniques. Rather than the sort of spatialisation strategies one might find in a virtual reality application, which often define (and parameterise) things in terms of precisely specified

[3]For the purposes of this article we can understand stereo diffusion simply as the practice of scaling and routing a stereo signal to one or more (usually pairs) of loudspeakers in performance. The various strategies and approaches taken, and their relative merits, is naturally a rather more complex topic.

https://doi.org/10.1017/S1355771810000312 Published online by Cambridge University Press

locations, a typical BEAST diffusion strategy might be more concerned with interpreting the character and behavioural implications of the stored sonic image (intimate, distant, high, low, dramatic, 'in your face', slowly evolving, erratic, etc.) and turning them into an appropriate acoustic reality in the performance space – enhancing the effect of distance, for example, through the use of speakers which really are distant from the audience, or a sense of intimacy, in which the sound seems to be whispering in the listener's ear, through deploying speakers as close to the listeners as possible. When limiting oneself to the requirements of strict localisation the latter might well be impossible to achieve in a large performance space with a reasonable sized audience, but a pragmatic approach allows for an elegant solution to this dilemma. If one ensures that there is a pair of speakers close to each section of the audience, and routes the signal to every such pair, each audience member will perceive the signal as intimate, locating it differently depending on which speakers are most proximate. It is worth noting that such an approach makes a virtue of precedence effect, usually the bane of spatialisation involving large audience groups crowding into a sweet spot. In fact precedence effect can play a large role in the success of stereo diffusion practice in general.[4]

We would like to propose what some may consider a controversial notion: although precise localisation and/or soundfield reproduction might be important for audiovisual or virtual reality applications (although the limitations of current cinematic presentation schemes suggest that a surprising amount of imprecision will be tolerated by audiences), for *musical* or artistic purposes it is less important than more general qualities, such as rough localisation based on qualitative considerations (e.g. front-ness, or a sense of height), dynamic aspects such as degrees and types of motion, degrees of diffuseness[5] versus localisation, non-specific spatial perceptions such as a sense of envelopment, and other less concrete considerations such as maintaining transparency, articulating contrapuntal aspects, and so on. Such spatialisations need not be conceived of in terms of real world spatial scenes, but could instead be realised using artistic (i.e. compositional) and musical criteria.

If one accepts that this principle holds true for at least a significant subset of multichannel practice, then adaptability and pragmatism seem worth retaining as potentially useful guiding principles as one ventures into the frontier areas of multichannel composition.

5. NON-SPECIFIC VS. SPECIFIC APPROACHES TO SPATIALISATION

The notions of adaptability and pragmatism can be seen as tangibly active in both past and recent BEAST practice through the incorporation of what we will refer to herein as *non-specific approaches to spatialisation*. These can be understood loosely as approaches that do not attempt to simulate precise locations and or directions of (usually point) sources. These can include the use of spatialisation as an abstract or artistic element, rather than to simulate a spatial scene understood in conventional terms based in real-world experience.[6]

Conceptually speaking, diffusion falls partly under this heading, at least when working with a touring system such as BEAST, since the precise speaker configuration and locations will generally not be known in advance (see figures 1–3 for similarities and differences in system designs for three different spaces in Birmingham). That said, it also embodies a certain amount of specificity, since a coherent stereo or multichannel image may be encoded within the work, since the directions of the loudspeakers may be somewhat specified (at least relative to a sweet spot), and because the diffuser may act with clear intent in terms of how the sound should be localised. Some non-specific approaches can be found in the discussion below.

6. NEW APPROACHES, TECHNIQUES AND RESOURCES

Below we will discuss some of the aspects of and resources for multichannel composition used with or developed for BEAST over recent years. As will become apparent, many of the developments discussed have resulted in a blurring of the line between composition and performance. Diffusion practice in some sense arguably does this already, at least insofar as one considers spatial diffusion as 'completing' a work.

While the question of the status of diffusion as performance is certainly interesting in its own right, within the BEAST community diffusion has largely been adopted and developed for its usefulness

[4]We should say, as a pre-emptive response to obvious objections, that we would be the last people to suggest that such approaches are appropriate to *every* sort of material. Precedence effect can work just as well against one's intentions, and as a general rule material which requires a clear and stable stereo image to be maintained will not fare well under such treatment – all of which is another way of saying that bad diffusion is as easy to find as bad performance.

[5]Note that the term diffusion is used herein to refer both to the performance practice, and to the qualitative aspect of sound which describes its relative localisability within a space. We trust that the active usage is made be clear by context.

[6]In some cases this may even go so far as to involve treating the loudspeakers 'instrumentally', for example as point sources or elements in a musical aggregate, which might perhaps result in one instance of what Denis Smalley refers to as 'technological listening' (Smalley 1997).

https://doi.org/10.1017/S1355771810000312 Published online by Cambridge University Press

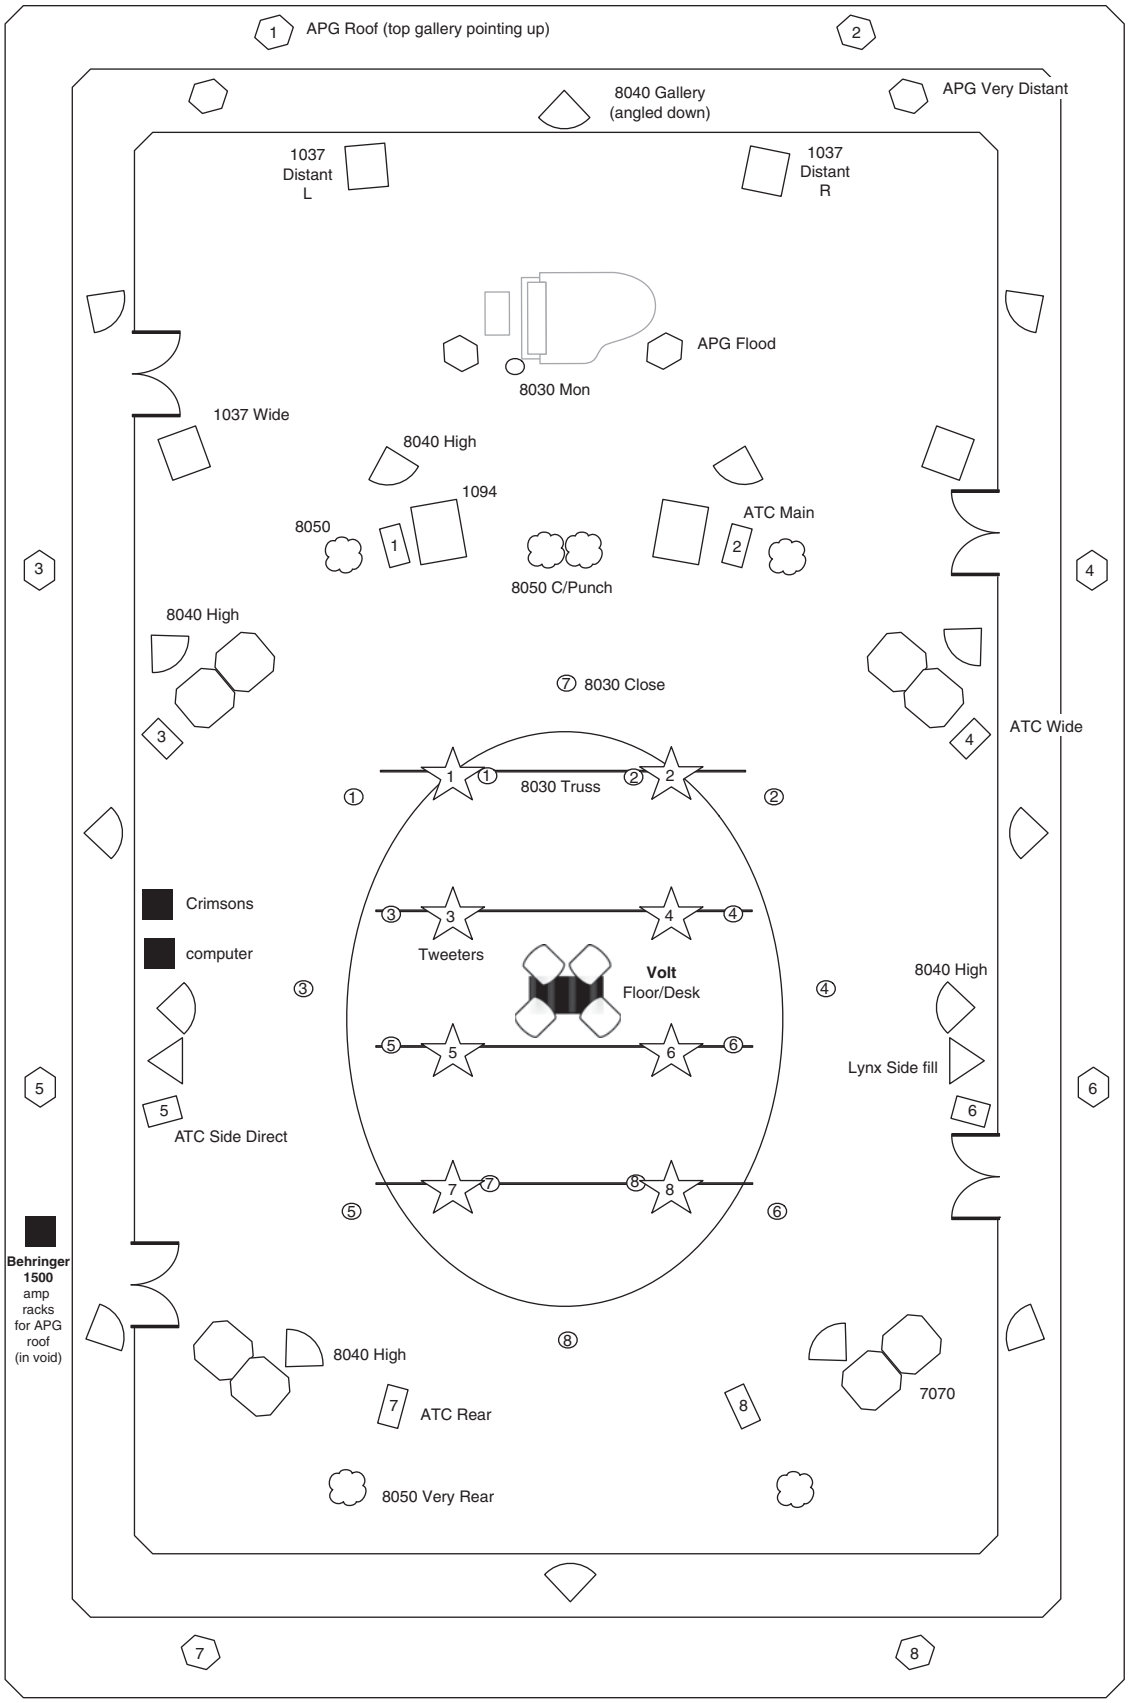


Figure 1. BEAST system layout at the CBSO Centre, Birmingham, May 2009.

https://doi.org/10.1017/S1355771810000312 Published online by Cambridge University Press

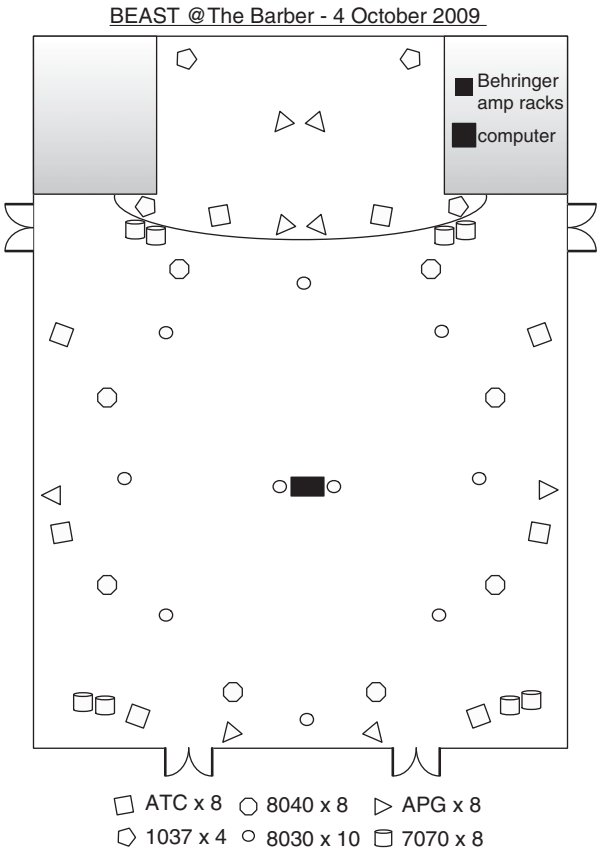


Figure 2. BEAST system layout at the Barber Institute, Birmingham, October 2009.

(or even necessity); that is, on pragmatic grounds rather than as part of a larger culture which valorises the diffuser qua performer. Given this, the distinction being discussed should be understood as between decisions which are taken or processes which are applied prior to the piece reaching its 'finished' state (something which – depending on one's way of working – may be difficult to define) and those which are taken after that, in order to adapt the piece to a given presentation context. Whether or not those later processes are performed or automated is not relevant in terms of the present discussion. The sections below discuss both 'pre-finished' and 'post-finished' approaches, and in many cases the techniques described could be used for either. Note that the sections are heterogeneous in nature and should not be understood as constituting strictly discrete or orthogonal categories.

6.1. BEASTmulch

The primary platform for this ongoing exploration into multichannel composition in recent years has been BEASTmulch,[7] an AHRC-funded research project whose primary outputs comprise two major elements: *BEASTmulchLib*, a SuperCollider[8] class library designed for use in the creation, processing and presentation of complex multichannel signal chains; and *BEASTmulch System*, a software application based on the library and designed for the presentation of multichannel audio works over complex loudspeaker arrays. The latter is the software component of the BEAST concert system.

In brief, BEASTmulch (figure 4) allows one to distribute arbitrary numbers of audio channels between inputs and outputs, to scale or process the signals individually or as groups at any stage of the chain, and to automate control of this. It includes support for a variety of well-known techniques, such as Vector Base Amplitude Panning (VBAP) (Pulkki 2001) and Ambisonics, as well as other more idiosyncratic approaches. It has support for a variety of controller devices, including MIDI controllers, the IRCAM EtherSense, GUI controllers, and our own custom built OSC fader boards.[9] The software includes data about speaker type and location, which some spatialisation approaches depend upon. Adaptability and pragmatism have been guiding principles in the software's design, not least because rehearsal opportunities are generally very limited.

6.2. Multichannel diffusion

A natural outgrowth of stereo diffusion practice has been multichannel diffusion, which was first done with Max/MSP and a simple MIDI fader box in a BEAST event as part of the *20/20 Re:Vision* twentieth anniversary weekend, 7–9 March 2003, at the CBSO Centre in Birmingham. In essence this approach is similar to its stereo version – that is, it is based on mixing between different sets of speakers in combination – but uses a source medium of greater than two channels. Most commonly this has been done with eight channel pieces, but it has also been done with other channel configurations, for example 5.1. Thus for an eight-channel piece one might have a close ring, a distant ring, a high and/or overhead ring, one or more 'special effects' arrays, and so forth.

This way of working is now well established within BEAST practice and makes an appearance in most BEAST events. Works conceived for this way of working, and which are based on a standard format such as an eight-channel ring or a 5.1 array, have the advantage of maintaining broad compatibility in a technical sense with any other systems designed around those standards, whether those systems offer opportunities for diffusion or not. (Artistically speaking of

[7]www.beast.bham.ac.uk/research/mulch.shtml.

[8]http://supercollider.sourceforge.net.

[9]Designed by Sukandar Kartadinata and based on his gluion interface: www.glui.de.

https://doi.org/10.1017/S1355771810000312 Published online by Cambridge University Press

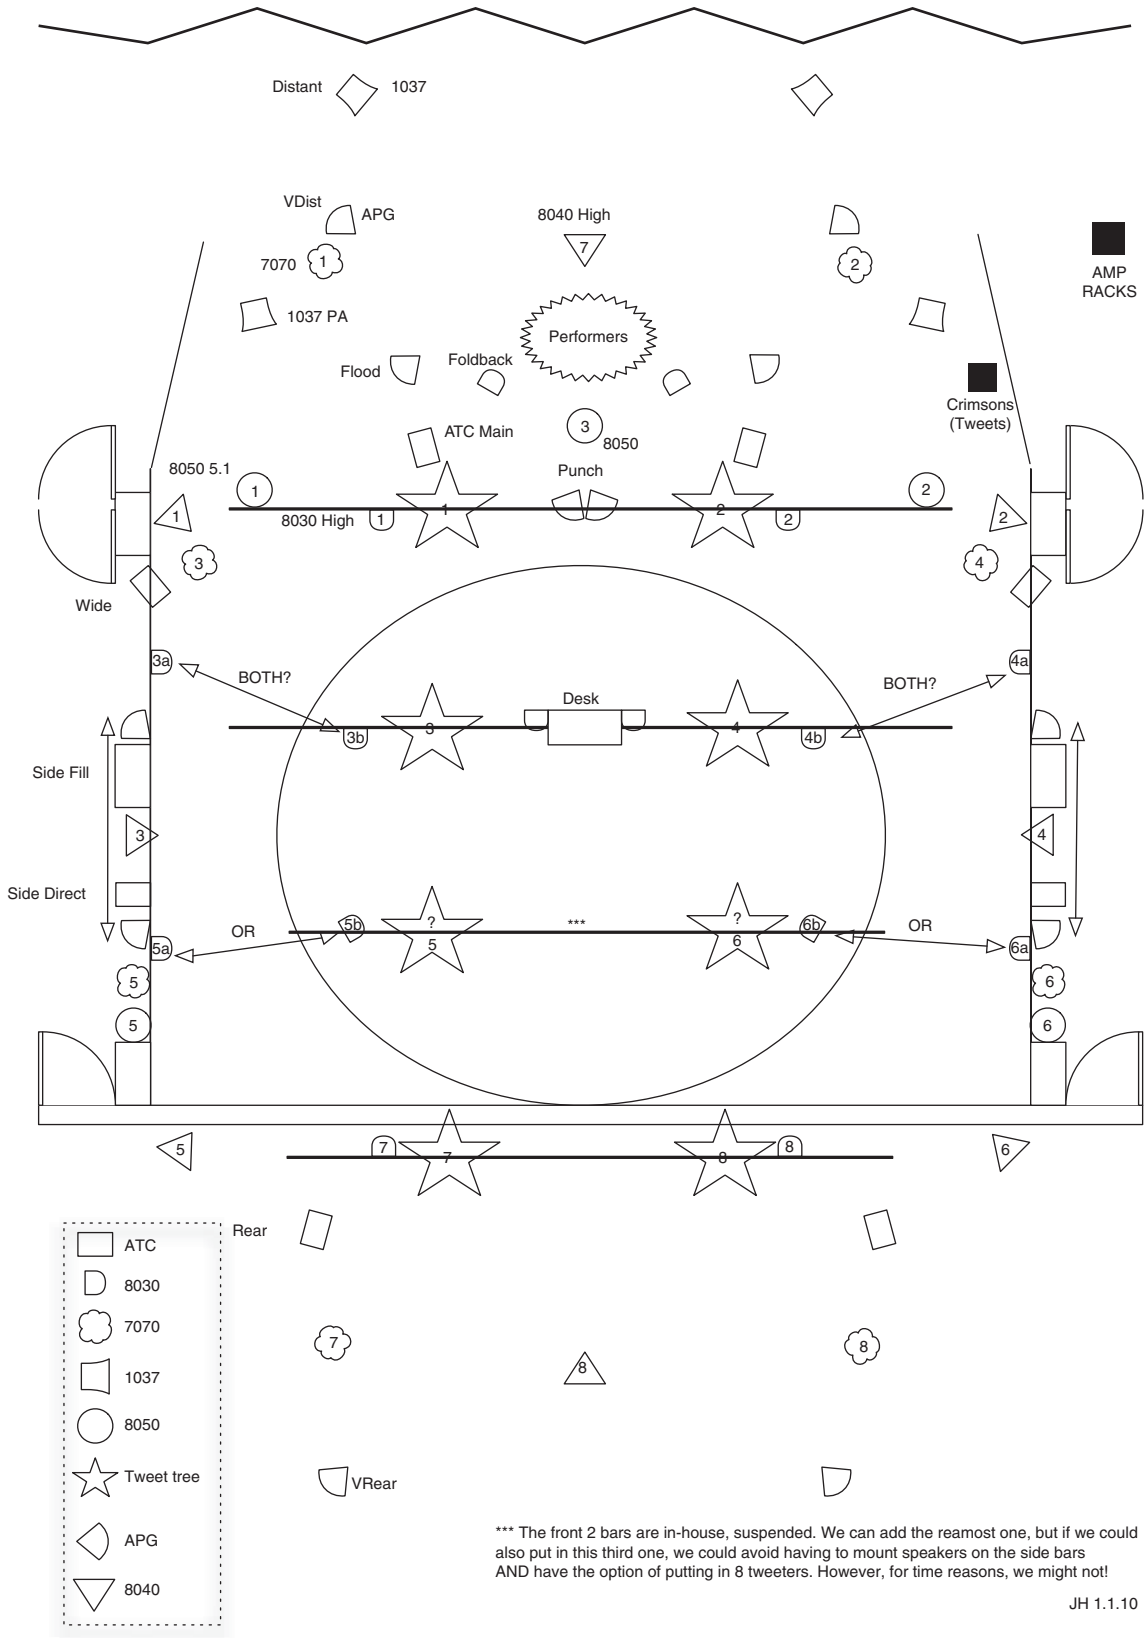


Figure 3. BEAST system layout at George Cadbury Hall, Birmingham, January 2010.

https://doi.org/10.1017/S1355771810000312 Published online by Cambridge University Press



Figure 4. A test setup for the BEASTmulch project.

course this is only true to the extent that the composer has conceived of the work's multichannel diffusion in a pragmatic fashion, i.e. considering what one would do with a high array, *providing one is available*, and keeping in mind alternative strategies for situations where one is not.)

6.3. Composing in stems

Multichannel diffusion is sensible when one wishes to work with channel configurations based on standards or pseudo-standards, but if one wishes to have larger numbers of discretely routable channels or sources, this becomes problematic given the arguable lack of any standard loudspeaker layouts with more than eight channels (even source channel to speaker numbering in 5.1 and eight-channel works is not consistent). Keeping pragmatism and adaptability in mind, a number of recent BEAST compositions have adapted an approach and term from mastering practice, the grouping of elements into *stems*. Stems constitute the submixes or – more generally speaking – discretely controllable elements which mastering engineers use to create their final mixes. In a similar fashion, one can compose in stems, separating out elements that need to be treated discretely in a final spatialisation, which in itself may vary to a small or great extent from one performance to another. In BEAST parlance, stems may be mono, stereo or multichannel. As a simple example, one could imagine a piece consisting of two eight-channel stems, one intended for ear-level localisation, and one intended for a higher location. In a large-scale multichannel system that contained appropriate arrays one could route these stems as desired. In a smaller setup consisting of only eight channels they could both be routed to the same array. Multichannel stems can be further reduced in size of course, through mixing and/or processing, and one could easily imagine how such a piece might be straightforwardly adapted for a quad or stereo system. The division of material into stems need not be based on spatial location, however. One might easily imagine a piece consisting of stems for foreground and background or figure and landscape elements, or one distinguishing between moving stems versus static ones.

Note that composing in stems does not in itself imply any particular final distribution or technique. Prototype stem approaches were used within notional eight-channel works by then BEAST composers David Berezan and Hasnizam Abdul Wahid (a stereo pair for normal manual diffusion plus six channels of fixed material to originate in speakers placed very close to the audience); in *Rock'n'Roll* (2004) Jonty Harrison somewhat inverted this approach by using a stereo track as a close, focused, central image and the remaining six as a diffuse, environmental image; both images were conceived as independently diffusible, should appropriate speakers be available. Another early example of stem-based composition with BEAST is Sergio Luque's *Happy Birthday* (2006), which consists of three stereo stems, intended for near, middle and far presentation, respectively. This work has been presented successfully in a number of realisations including a stereo mixdown, using the Game of Life Wavefield Synthesis System in the Netherlands, on a ten-channel system as part of the Integra Festival, and with each stem treated to individual stereo diffusion over appropriate subsets of the BEAST system.

Since BEASTmulch supports arbitrary numbers of inputs and outputs one can easily make decisions about stem treatment at the time of rehearsal, but one could just as well produce mixed-down versions for a given configuration. For those that might feel that a stems-based approach could prove onerous, it is worth noting that one need only create a version for a given configuration (such as 5.1) once. Those who are forced to render out their stems to more than two or three such variations might just as easily count themselves lucky to have so many opportunities, as feel burdened by the requirement!

Similarly, while composing with stems might lack some of the appeal of composing to encoded formats such as B-format and DirAC (Pulkki 2007) (which are essentially *specific* in conception if not always in realisation), wherein decoding to a given setup may be largely or completely automatable, stems as an intermediate format retain the advantage of allowing pragmatic adaptation of the work's spatialisation in a fashion which makes the best (possibly *non-specific*) uses of the resources at hand (i.e. the nature of the system or performance venue).

Naturally this approach blurs the boundaries between composition and performance, since some choices (or at least their specific realisation) are

https://doi.org/10.1017/S1355771810000312 Published online by Cambridge University Press

deferred at least until the performance situation is known. The 'finished' work is necessarily in some sense not quite finished. This is arguably even truer with stem-based composition than is the case with stereo diffusion.[10]

6.4. n-channel composition

This approach is both related to and overlaps with stem-based composition, and can take stems as source materials or produce them as outputs, but also includes the possibility of inputs and outputs of dynamic sizes. Broadly speaking n-channel composition is an approach in which sources may be spatialised over arbitrary sub-arrays whose sizes are not known in advance. The output 'stems' themselves need not have fixed numbers of channels, but might be varied according to what is available, and possibly generated in real-time. These outputs can be hard assigned or diffused as desired, taking advantage of the available resources or compensating for deficiencies in the concert situation that could not have been anticipated at an earlier stage in the compositional process.

n-channel approaches again blur the boundaries of composition and performance, although in most cases the specific techniques described can be utilised in both non-real-time and real-time (i.e. in performance) fashion. (Naturally the available computing resources do still sometimes impose practical limitations.) Whether used for real-time realisation or not, however, it is the practice of adapting to specific performance circumstances (i.e. system, space, etc.) that is the essential quality of n-channel approaches.

We will discuss a number of techniques that have been used for n-channel composition in BEAST works below, which while not necessarily limited to it, are nevertheless well suited to this strategy.

6.4.1. Vector Base Amplitude Panning

BEASTmulchLib contains a SuperCollider port of Ville Pulkki's Vector Base Amplitude Panning (VBAP) (Pulkki 2001). Simply put, VBAP allows one to equal-power pan sources over arbitrarily spaced 2D and 3D arrays of speakers. The algorithm makes certain assumptions, most importantly that the speakers are all equally distant from a central point; that is, that you are dealing with a complete or partial ring or sphere. This is naturally important because of the precedence effect, but with a system such as BEAST – wherein speakers usually play multiple roles within an event – this is not always possible. One can, however, compensate to some extent using delays on the signals sent to the closer speakers. Control parameters for VBAP take the form of azimuth and elevation values expressed in degrees. Signals are generally panned between pairs (in a 2D ring) or triplets (in a 3D dome or sphere) of speakers. When dealing with a reasonably large number of speakers (so that the angles between them are less than the 60° stereo standard) this approach makes localisation of 'point source' material relatively robust; in other words, even listeners somewhat removed from the centre of the sweet spot will locate the panned sound in roughly the right direction. Thus although VBAP is certainly *specific* in its conception of spatialisation, it does fare well under situations where somewhat *non-specific* spatialisation is an inevitable result, such as concert situations where much of the audience's location relative to loudspeaker arrays strains the definition of 'sweet'.

What makes VBAP a useful 'n-channel' technique is that sound direction is abstracted from speaker location. The same control parameters can be used with a variety of loudspeaker configurations. As noted above, the number and spacing of the loudspeakers along the ring or sphere is in principle arbitrary and can be uneven, although naturally some configurations will perform better than others, and it is not possible to pan effectively across large gaps.

6.4.2. Spatial decorrelation techniques

For the purposes of this article we will consider spatial decorrelation techniques as including any approach which produces usefully decorrelated versions of a signal across two or more channels. For a more detailed discussion of the theory behind such techniques and of one approach see (Kendall 1995), but the important thing is that the resultant signals are decorrelated in a manner which produces spatialisations with enhanced diffusion, and the concomitant qualities of increased volume perception, envelopment and so on, depending of course on the individual technique and the parameters used. Normally one would limit the definition to those approaches which are relatively artefact free, but in our discussion of such practices within the BEAST community we also include under this heading those which significantly alter the source material (perhaps resulting in decorrelation only as a side effect rather than the composer's primary intention).

Certain types of multichannel granulation fall into this latter category, for instance. These include approaches which allow for granulation across selected sub-arrays of the BEAST system, allowing for localisation with a variation in physical volume perception, or granulating the same source in variants simultaneously across multiple sub-arrays, for example

[10]It seems reasonable to assume that most composers of works intended for stereo diffusion consider the stereo source tracks themselves an acceptable version of the piece, albeit perhaps not one that is ideal for all listening situations.

https://doi.org/10.1017/S1355771810000312 Published online by Cambridge University Press

changing the pitch for sub-arrays at different heights. Such an approach is an example of an artistically motivated *non-specific* approach, since it is primarily about spatial registration and transparency rather than the simulation of a spatial scene understood in real-world terms.

Decorrelation effects in general are relatively robust across a variety of audience positions, and normally one must be fairly close to one loudspeaker (i.e. close enough to bring masking into operation) before issues arise. A group of decorrelated signals (whether in one locale or surrounding the audience) will thus tend to maintain its relative locale (or lack thereof) across a surprisingly large sweet spot (thus making such approaches eminently pragmatic in a concert context). Indeed since precedence effect is not largely at play with these, they are well suited for use with systems such as BEAST, which contain non-homogenous arrays of loudspeakers at different *distances*. One notable granular variant which takes advantage of this possibility to simulate physical depth is our Spatial Swarm Granulation technique, again developed as part of the BEASTmulch project (Wilson 2008). This approach uses an adapted boids algorithm (Reynolds 1987) to model grain position, and a nearest neighbour approach for output mapping. Other similar granular techniques which allow movement of granular output across distributed non-homogenous arrays are under development.

Other types of spatial decorrelation techniques used with BEAST include FFT-based approaches such as those described in Kendall (1995), Topper, Burtner and Serafin (2002) and Torchia and Lippe (2004). A number of variants have been developed, but most involve either the decorrelation of phase in two or more altered copies of a source signal, or the splitting of each bin's magnitude to a number of outputs.

6.4.3. Spectral splitting

One notable approach related to spatial decorrelation is 'spectral splitting'. This notion perhaps has it origins in non-homogenous loudspeaker systems intended for diffusion, in which, due to the varying frequency responses of different loudspeakers, their relative proximity and orientation, the onset times of different sounds or components, and a number of psychoacoustic considerations, sounds seem to separate out spatially to different parts of the array. In the BEAST system this approach can perhaps most easily be seen in the use of specialised tweeter trees. These are generally suspended above the audience, and thus in many cases are the closest active speakers. The splitting effect is enhanced through the use of high pass filters (usually at about 10kHz) to reduce the frequency content of the tweeter's input signals. Again this is an example of making a virtue of precedence effect.

One (perhaps surprising) complaint that has been made about the current BEAST system is that the different loudspeaker models in use are too homogenous and consistent, and lack enough 'colour' to really bring spectral splitting effects to life when diffusing (it is true that, following a grant from the UK government's Science Research and Investment Fund in 2004–05, BEAST bought loudspeakers in matched sets of eight, rather than in stereo pairs, as the move towards 'eight-channel as the new stereo' was already clear). As BEAST concert programmes now standardly include eight-channel works, and multichannel diffusion is an established practice, it is desirable that multiple eight-channel arrays be available. Interestingly, the complaint of lack of colour has come mostly from composers diffusing stereo works, but it should be pointed out that it may well not be necessary, and is certainly not obligatory, to use all 80+ speakers to deliver a good stereo diffusion – it is perfectly possible to use just two or four speakers from any given eight-channel array. One good example of speakers in an eight-channel array not being the most appropriate for stereo use is the distinction between what are referred to as 'Side Direct' speakers, often essential as two speakers in a ring for eight-channel works, and 'Side Fill' speakers (Harrison 1998, 2000) which are generally off-axis; that is, placed for diffuse effect, for example pointing at walls or out of direct sight. The main function of Side Fill speakers in stereo works may be to provide a smooth, subtle link when crossfading between the front and rear speaker arrays, while Side Directs, on the other hand, may be deemed to be too much like wearing a giant pair of headphones to be of much use in a stereo context. BEAST normally supplies both, as the attached system schematics show (see figures 1–3. Of course, even with the provision of Side Fills, the Side Directs can always be called into use as a 'special effect' in stereo diffusion).

In any case, 'colour' can of course be simulated, and this has been done with BEAST; for example by filtering the signals routed to some off-axis speakers at the very front of the hall in order to enhance an effect of distance. On a more elaborate scale, Garfield Benjamin's *In the Eye* (2009) is a work designed specifically with this approach in mind, and attempts to create 'auto-diffusion' effects with frequency variant multichannel signals which are routed to a number of sub-arrays, each of which is treated to different types of filtering.

Scott Wilson has also developed a number of idiosyncratic additive spectral analysis-resynthesis approaches to multichannel spatialisation that could be described under this general heading (as well as under decorrelation approaches above, depending on usage) using Kelly Fitz and Lippold Haken's Loris library (Fitz, Haken, Lefvert and O'Donnell 2002).

https://doi.org/10.1017/S1355771810000312 Published online by Cambridge University Press

6.4.4. An example of n-channel composition

An example of an n-channel work is Wilson's *Gotlandic Miscellanea* (2008). This work is based on a series of mid-side format stereo recordings, which are processed and spatialised in real-time over arbitrarily sized speaker arrays using techniques such as those described above. It also makes use of a granulation approach which determines routing based on input amplitude, increasing the perceived physical volume of the result by expanding the number of outputs from the front to the back of the space as the amplitude increases. The arrays selected for this in a given performance context are chosen based on qualitative and pragmatic factors; for example trying to maximise a sense of 'sweep', or envelopment, given the speakers available. With a small system (e.g. a ring of eight speakers), there may be little or no choice, but with a large system, there may be several good alternatives. This amplitude-dependent approach works in a similar fashion to what is arguably 'good' stereo diffusion practice (i.e. spatialising 'with' a musical gesture rather than 'against' it), but by automating the process allows for a degree of detail and control that might be difficult to achieve manually in performance. Other amplitude variable spatialisation approaches are used in the piece as well, for instance to control image width.

The work also demonstrates a strategy which is useful across a variety of approaches: in most cases the left–right relationship of the stereo recordings is maintained: for example, when spectral magnitude-based decorrelation is used, the left signal is decorrelated only to loudspeakers on the left and the right signal only to those on the right. This helps to maintain a sense of the encoded stereo image, despite it being decorrelated across eight or more spatially separated speakers, and has proved very useful in particular when attempting to spatialise environmental recordings in a manner combining envelopment with localisation of particular aspects.

The work has proven extremely (and relatively straightforwardly) adaptable, having been premiered as part of the *Integra Festival* in 2008 in Birmingham over a ten-channel system, and subsequently presented on two large-scale BEAST setups of 60+ channels. An eight-channel version has been created, and a 5.1 version is planned. Finally, a stereo version was released as part of the *Deep Wireless 6* CD compilation.

6.5. Hybrid approaches

Hybrid approaches are naturally also possible, and in truth most of the works mentioned above can be understood as hybridised in some sense. Eric Bumstead's *|kroʊn|* (2007) and *BlckWnd* (2009) are works which contain both an eight-channel stem – normally hard-assigned to a single array in concert – and a stereo stem intended for stereo diffusion (perhaps in part over the same array as the eight-channel stem; a combination which BEASTmulch System supports). The former work also contains two monophonic 'special' stems: a 'vocal' track, intended to be assigned to a central position in front of the audience, and an 'as distant as possible' track. Both of these can be treated as appropriate (perhaps with real-time enhancement) given the available resources.

Zlatko Baracskai's *Culpable Passage* (2008) consists of three eight-channel stems, conceived of in musical terms as foreground, moving background, and steady background. Each stem could be diffused to a number of sub-arrays, and may also be treated to a number of real-time processes such as delays and rotations. It thus combines diffusion, stems, n-channel realisation, live processing, and aspects of specificity and non-specificity.

7. CONCLUSIONS AND CAVEATS

One thing which we have not done in the text above is make detailed qualitative assessments of the approaches outlined. While the BEASTmulch project and our related activities, such as giving concerts, have afforded us the opportunity to test and refine these approaches using a large-scale system, the sort of formal study or survey which would be required to make qualitative comparisons in a detailed manner was not within its remit. (It seems likely that such assessments would be highly context-dependent in any case.) That said, all of the approaches described above have seen use in concert practice, and were deemed (at least on subjective grounds) to be successful enough to warrant further use and exploration.

While it is easy to be excited by the opportunities that such developments represent, it is important to keep in mind that at the same time they bring with them the possibility of negative side-effects; for example the breakdown of standards such as the 'eight-channel ring piece', and the loss of the easy exchange of artistic works (and thus to some extent also of the aesthetic and technical knowledge which they embody). Perhaps more worrying is the possibility of a partial reversal of the democratising effects of cheap multichannel audio and computer hardware through the growth of institutionally affiliated large-scale multichannel presentation and research systems – such as the ZKM Klangdom (Ramakrishnan, Goßmann and Brümmer 2006), the Sonic Lab at the Sonic Arts Research Centre in Belfast, the Allosphere at the University of California in Santa Barbara (Amatriain, Höllerer, Kuchera-Morin and Pope 2007), and of course BEAST itself – with all the issues of access and exclusivity that one associates with the early institutionally based days of electroacoustic music history. While not a remedy for this situation,

https://doi.org/10.1017/S1355771810000312 Published online by Cambridge University Press

some of the pragmatic approaches discussed above at least allow for composers to work in a way which will allow them to fairly easily adapt their works to such systems when opportunities arise, as well as allowing for works composed for those (often rather idiosyncratic) systems to have a life outside of the institutions in which they are based.

Also of concern are the barriers to entry presented by technical considerations, and the question of what knowledge and skills can or should be assumed when designing software systems. One form this has taken in the BEASTmulch project has been the question of where to draw the line between using the library and using the system application. Put another way, what sorts of things are general enough to be worth including in the application, and what sort of things will ultimately require programming skills. Many of the pieces described above have been presented using BEASTmulch System, but others required custom implementations (using BEASTmulchLib and SuperCollider or other tools) to realise. (In essence these works are themselves standalone pieces of software.) Although one can always make a given piece of functionality more accessible by adding an interface which does not require programming skills (a process which is certainly ongoing with BEASTmulch), it seems obvious that techniques which will require new or custom programming will continue to arise. Things like the amplitude-controlled spatial granulation described above, and in general new spatialisation approaches along the lines of 'Adaptive Digital Audio Effects' (Verfaille and Arfib 2001) are just one example of such an area. Given this, it is important that spatialisation environments provide hooks for expansion, allowing users to extend the system whilst still being able to leverage existing functionality. One BEASTmulch line of research in this direction is currently moving towards a new type of timeline GUI element[11] which will allow users to embed custom SuperCollider code within BEASTmulch System, and control its execution flow in a typical DAW-like fashion. An interesting side-effect of this development, related to the blurring of composition and performance discussed above, is that BEASTmulch System may move from being primarily a performance environment to also being a potential composition environment in its own right. Whether this approach will prove tenable in the long run remains to be seen. Another proposed line of development is towards finding ways to integrate the system with typical existing composer workflows. The risk of this direction is that it may lead to inelegant or 'kludgey' solutions, but the *pragmatism* involved and the potential improved accessibility it might offer are definitely appealing.

[11]Inspired in part by Jan Trützschler's Tea Tracks, but not implemented using it.

Another issue raised is the question of storage and exchange formats. While an approach such as VBAP is convenient in that all it requires at the time of composition (as opposed to performance) is a source signal and some control data, what is not convenient is that there are no standard formats for storing material and control data in a VBAP-friendly way. Non-specific approaches present even more problems in this regard, since the qualitative descriptors involved, while certainly meaningful and useful, may not lend themselves to automated interpretation. Recent discussions around the development of a SpatDIF file format, which would allow for straightforward interchange of spatial information (Kendall, Peters, Geier, Telekom and Berlin 2008), hold some promise however, although the issue of how to describe non-specific qualities, which seem to some extent dependent on subjective language, remains difficult. This concern of storage and transmission is necessarily tied up with a blurring of the boundaries between composition and presentation, and presents practical archival concerns for the future as well as portability problems in the present. In any case a better storage format is not a panacea for the issues raised by the rise of large-scale multichannel composition: while such a development might swing practice back towards the composition side of the divide, it seems likely that some aspects of adaptation will always require a certain amount of 'hand-tuning'.

The field of large-scale multichannel composition remains in its infancy, and despite these concerns it is still easy to look with a certain amount of optimism at the possibilities which are emerging. Issues of access are improving as more systems come into being, and while these issues are likely to remain a concern, they are perhaps somewhat akin to those faced by the aspiring orchestral composer, albeit arguably with less (or at least less consistent and monolithically defined) cultural history getting in the way of what might be possible.

REFERENCES

Amatriain, X., Höllerer, T., Kuchera-Morin, J. and Pope, S. 2007. Immersive Audio and Music in the Allosphere. *International Computer Music Conference*, Copenhagen, Denmark. Ann Arbor, MI: Scholarly Publishing Office, University of Michigan: 276–83.

Fitz, K., Haken, L., Lefvert, S. and O'Donnell, M. 2002. Sound Morphing Using Loris and the Reassigned Bandwdith-Enhanced Additive Sound Model: Practice and Applications. *Computer Music Conference*, Gothenburg, Sweden. Ann Arbor, MI: Scholarly Publishing Office, University of Michigan: 393–400.

Harrison, J. 1988. Space and the BEAST Concert Diffusion System. In Francis Dhomont (ed.), *L'espace du son*. Ohain, Belgium: Musiques et Recherches.

https://doi.org/10.1017/S1355771810000312 Published online by Cambridge University Press

Harrison, J. 1998. Sound, Space, Sculpture: Some Thoughts on the 'What', 'How' and 'Why'of Sound Diffusion. *Organised Sound* **3**(2): 177–26.

Harrison, J. 2000. Diffusion: Theories and Practices, with Particular Reference to the BEAST System. *eContact* 2.4 on the website of the Canadian Electroacoustic Community/ Communauté électroacoustique canadienne: http://cec.concordia.ca/econtact/Diffusion/diffindex.htm.

Harrison, J. 2004. *Rock'n'Roll*. On *Environs*, IMED0788, *empreintes* DIGITAL*es*, Montreal, QC, 2007.

Kendall, G. 1995. The Decorrelation of Audio Signals and its Impact on Spatial Imagery. *Computer Music Journal* **19**(4): 71–87.

Kendall, G., Peters, N., Geier, M., Telekom, D. and Berlin, G. 2008. Towards an Interchange Format for Spatial Audio Scenes. *Proceedings of the 2008 International Computer Music Conference*, Belfast, UK. Ann Arbor, MI: Scholarly Publishing Office, University of Michigan.

Pulkki, V. 2001. *Spatial Sound Generation and Perception by Amplitude Panning Techniques*. ScD dissertation, Sibelius Academy, Helsinki.

Pulkki, V. 2007. Applications of Directional Audio Coding in Audio. *Proceedings of the 19th International Congress of Acoustics*, Madrid, Spain.

Ramakrishnan, C., Goßmann, J. and Brümmer, L. 2006. The ZKM Klangdom. *Proceedings of the 2006 Conference on New Interfaces for Musical Expression*, Paris: IRCAM: 140–3.

Reynolds, C. 1987. Flocks, Herds and Schools: A Distributed Behavioral Model. *Proceedings of the 14th Annual Conference on Computer Graphics*. New York: ACM.

Smalley, D. 1997. Spectromorphology: Explaining Sound-Shapes. *Organised Sound* **2**(2): 107–26.

Topper, D., Burtner, M. and Serafin, S. 2002. Spatio-Operational Spectral (SOS) Synthesis. *Proceedings of the 5th International Conference on Digital Audio Effects (DAFX-02)*, Hamburg, Germany.

Torchia, R. and Lippe, C. 2004. Techniques for Multi-Channel Real-Time Spatial Distribution Using Frequency-Domain Processing. *Proceedings of the 2004 Conference on New Interfaces for Musical Expression (NIME04)*, Hamamatsu, Japan: 116–19.

Verfaille, V. and Arfib, D. 2001. A-DAFx: Adaptive Digital Audio Effects. *Proceedings of the COST G-6 Conference on Digital Audio Effects (DAFX-01)*, Limerick, Ireland.

Wilson, S. 2008. Spatial Swarm Granulation. *Proceedings of the 2008 International Computer Music Conference*, Belfast, UK: 1–4.

https://doi.org/10.1017/S1355771810000312 Published online by Cambridge University Press