The University of Sheffield

James R. Mooney

Sound Diffusion Systems for the Live Performance of Electroacoustic Music

An Inclusive Approach led by Technological and Aesthetical Consideration of the Electroacoustic Idiom and an Evaluation of Existing Systems

Volume 2

Chapters 4 – 6, Appendices and Bibliography

Submitted for the Degree of Ph.D. in the Faculty of Arts

Department of Music

August 2005

4. Sound Diffusion Systems

4.1. Introduction

In Chapters 1 and 2 an approach was made towards an understanding of the technological and aesthetic scopes, respectively, of the electroacoustic idiom, with a broad focus on the demands of the public performance scenario from each of these perspectives. In Chapter 3, some of the ways in which these issues are carried forward as contrasting approaches to the practice of sound diffusion itself were outlined. From a technological perspective, this has been explained in terms of the relationship between coherent audio source sets and coherent loudspeaker sets. In aesthetic terms, two distinct methodologies have been observed. Of course, these differing top-down and bottom-up aesthetics, as they have been defined, will have a fundamental impact upon the ways in which the technologies are approached in both composition and performance. From these three previous chapters an overall impression of the 'range' of electroacoustic music has been given. It has been proposed that if any specific sound diffusion system is to be of *widespread* use, then it will have to be able to accommodate works – and attitudes, and methodologies, and so on – from across this range side-by-side, in the context of a live performance.

The purpose of the present chapter is to evaluate the extent to which various existing sound diffusion systems are able to satisfy this challenging demand, and more importantly to examine the ways in which this is (or is not) achieved. As far as the author is aware, there is no single published document that collates, summarises, and evaluates current sound diffusion technology in this way. Observations made in this respect will form the basis of a proposed 'way forward' in the development of future systems; this will be given in Chapter 6. In order to facilitate the realisation of this overall objective more easily, a system of criteria for the evaluation of sound diffusion systems will be proposed. These criteria will also be used for the evaluation of the M2 Diffusion System (Chapter 5), which was co-designed, co-developed, and co-implemented by the author as an integral part of this research. The review process will begin with a non-specific explanation, in abstract terms, of what a diffusion system is.

4.2. What is a Sound Diffusion System?

A sound diffusion system is hardware, software, or a combination thereof, used to facilitate the presentation of electroacoustic works via loudspeakers in a live concert situation.

As such, any given diffusion system is likely to be subjected to a fairly weighty set of technological and aesthetic, not to mention practical, demands. From a technological perspective, the role of a sound diffusion system is to mediate between audio source(s) and a loudspeaker array consisting of multiple loudspeakers, via some kind of intermediate control interface that allows the performer to execute the diffusion as required. The format of the mediation is encoded audio streams (see section 1.6). An abstract graphical representation of a generalised diffusion system architecture is given in Figure 20, below.

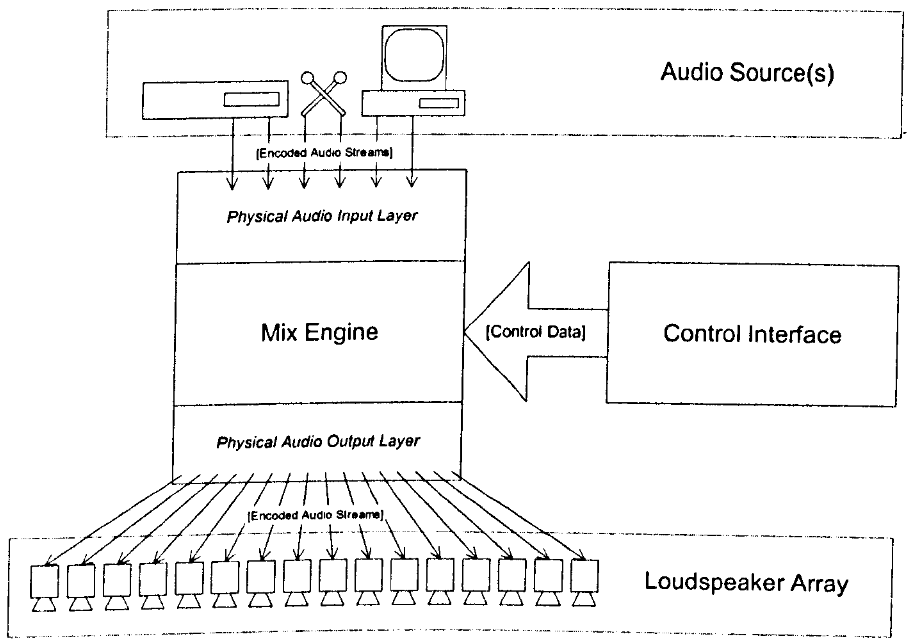

Figure 20. Graphical representation of a sound diffusion system in terms of four main components: audio source(s), control interface, mix engine, and loudspeaker array.

The point of entry for encoded audio source streams into the system is via a physical audio input layer, which receives encoded audio streams from the various audio sources (CD player, microphones, computer sound card outputs, *et cetera*; any device capable of outputting at least one transitorily encoded audio stream can potentially be used as an audio source). The transmission of encoded audio streams to loudspeakers (the point of exit) takes place via a physical audio output layer. These input and output layers are also illustrated in Figure 20. The mix engine mediates between these two layers, performing the necessary signal routing, mixing, and any other signal processing tasks. Often (but not necessarily always) this will involve outputting a number of encoded audio streams that is larger than the number of input streams. The mix engine performs its task according to control data received from the control interface, which is the means by which the performer interacts with the diffusion system. The input and output layers are necessarily physical, whereas the mix engine and control interface can be implemented in software or hardware or a combination of both.

The exact technological means by which the abstract process of sound diffusion is facilitated are various (a number of different possibilities will be described later in the chapter) but in many cases a hardware audio mixing desk serves as the physical audio input layer, control interface (faders), mix engine, and physical audio output layer. For any single performance, the loudspeaker array is most likely to be static, that is, it will consist of a fixed number of loudspeakers arranged in a set formation (although it is not unknown, in the author's experience, for loudspeakers to be added or removed, or their positions changed, within a single performance). The nature and number of channels contained within the audio source(s), however, is much more likely to change within a single performance. As a simple example, a concert programme may contain stereophonic fixed-medium-only works, multichannel fixed medium works, and works for instruments and tape. Consequently, the way in which the total number of encoded audio inputs is subdivided into coherent audio source sets may vary from work to work, and indeed the total number of audio inputs in use at any given time may also change. This is also likely to mean that the way in which the loudspeaker array is utilised – that is, subdivided into coherent loudspeaker sets – is also likely to change from piece to piece, even though the total number and formation of loudspeakers within the array remains constant. For example, a stereophonic work (representing a single two-channel CASS) diffused via an array of eight loudspeakers is likely to treat the loudspeaker array as four stereophonic CLSs; an octaphonic fixed medium work (if eight source channels have been treated collectively as a single CASS) is more likely to treat the entire loudspeaker array as a single octaphonic CLS.

Having briefly described the concept of the sound diffusion system in abstract technological terms, it can now also be said that a sound diffusion system will have to facilitate the presentation of electroacoustic works according to both top-down and bottom-up methodologies, as it is likely that both kinds of work will co-habit concert programmes. The following sections describe a set of criteria by which, it is proposed, sound diffusion systems can be evaluated in terms of their flexibility and appropriateness for the task of performing live electroacoustic music.

4.3. Criteria for the Evaluation of Sound Diffusion Systems

The objective of sound diffusion is, in a very general sense, to facilitate the presentation of one or more coherent audio source sets – via one or more coherent loudspeaker sets – such that the coherency of the audio source sets is maintained as far as possible. If this can be achieved then the diffuser will have done everything within his or her power to ensure that the discourse of the electroacoustic work will be effectively and faithfully communicated to members of the audience, and this – it is suggested – should be regarded as the ultimate aim of sound diffusion. It is clear that if this overall objective is to be uniformly achieved – given the broad technical and aesthetic scope of the electroacoustic idiom – a sound diffusion system will have to satisfy a complex set of interrelated demands. Some of these demands stem directly from the nature of electroacoustic music itself, in technological and aesthetic terms. In this particular respect, it is proposed that a sound diffusion system can be evaluated in terms of its:

4.3.1. Ability to Cater for the Technical Demands of Live Electroacoustic Music

The various technical demands that might be exerted by any given piece of electroacoustic music were outlined in sections 1.7 and 1.8, with more general observations with regard to how the technologies are appropriated by electroacoustic musicians being made in the latter sections of Chapter 1. In the context of the diffusion system, ideally, all of the possible permutations should be easily manageable from a technical standpoint (it will later be argued that this is not currently the case). It should, of course, be noted that the technological repertoire is constantly expanding.

As an arbitrary example, this means that an ideal sound diffusion system should be able to accommodate the technology required to perform an electroacoustic work for live instruments, live synthesis, and tape, alongside that required for fixed-medium-only works, and alongside that demanded by works for instruments and 'live electronics.' It must be remembered that a concert programme of electroacoustic music is (at least, in the author's experience) reasonably likely to incorporate works of *varying* technical demands in this way. The diffusion hardware (or software if applicable) must therefore be designed in such a way that these demands can be met without unduly disrupting the flow of the programme as a whole. Hiatus between consecutive works within the concert programme should be avoided if at all possible; this can arise if works of varying technical demands necessitate the physical re-patching of audio hardware, for example.

In section 1.6.2 it was noted that individual encoded audio streams are often used collectively to abstractly represent spatio-auditory attributes; in section 3.5 the coherent audio source set (CASS) was presented as a convenient conceptual framework for representing this. In the present context it should be noted that, in addition to those technological demands already discussed, electroacoustic works are *also* likely to vary in terms of their CASS demands. The number and individual nature of CASSs in any given electroacoustic work is, at the general level, an unknown variable. A stereophonic fixed medium work is most likely to consist of a single, two-channel, CASS: this consensus is practically unanimous and therefore relatively simple to deal with. When the number of CASSs (and the number of constituent channels within any given CASS) increases, the situation becomes exponentially more difficult to accommodate from a technological standpoint, particularly given the relatively inflexible signal routing offered by most studio mixing desks. If physical re-patching is to be avoided, and if works comprising variable numbers of coherent audio source sets – themselves consisting of a potentially arbitrary number of constituent channels – are to be accommodated, then a high degree of dynamic signal routing flexibility is clearly required, particularly if works that vary in these respects are to be performed consecutively; this will be discussed more fully in section 4.3.3.

4.3.2. Ability to Cater for both Top-Down and Bottom-Up Approaches to Sound Diffusion

In sections 3.12 to 3.14, some of the characteristics that differentiate top-down and bottom-up approaches to sound diffusion were proposed and summarised. As concert programmes are fairly likely to incorporate works conceived from both of these contrasting standpoints, it is desirable that diffusion systems should be able to cater for them equally and without prejudice. It has been noted, for instance, that top-down diffusers tend to favour pluriphonic relationships between a single CASS and multiple CLSs, while bottom-up diffusers are more inclined to adopt a uniphonic approach whereby each channel of the CASS is represented by a one loudspeaker within a single CLS. The former (pluriphonic) approach is reasonably flexible because appropriate CLSs can be defined as necessary within a variety of different loudspeaker arrays. The uniphonic approach can be more restrictive as it usually relies on a fairly specific set of loudspeakers in terms of relative positioning. If both models are to be supported, this means that the loudspeaker array must allow appropriate CLSs to be defined in each case. The array shown in Figure 21, for example, would allow for both stereophonic and 5.1 CLSs to be easily defined within it. A stereophonic work could be diffused pluriphonically by a top-down diffuser (see Figure 21(b)), while a 5.1 work could equally be presented uniphonically by a bottom-up diffuser (Figure 21(a)).

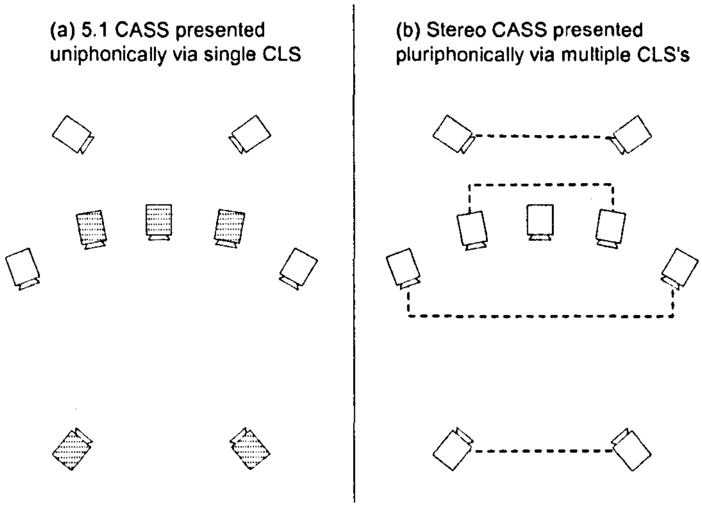

Figure 21. Loudspeaker array allowing for the uniphonic presentation of a 5.1 CASS in addition to the pluriphonic presentation of a stereophonic CASS.

Of course if the scenario illustrated in Figure 21 is to be facilitated in practice, this raises further technological issues, particularly with respect to signal routing and mixing (to be discussed in section 4.3.3) and interface ergonomics (section 4.3.4).

It has also been noted that top-down diffusers are more inclined to favour loudspeakers that 'colour' the sound, while bottom-up diffusers tend to prefer loudspeakers that are 'transparent' in terms of their frequency response. Again, if both approaches are to be supported within the context of a single performance, then some kind of compromise is required. If the loudspeaker array consists of transparent loudspeakers, then perhaps some facility to add 'colouring' during performance is necessary. Alternatively, if the array consists in 'coloured' loudspeakers, then the ability to *compensate* for this during performance would seem desirable. In either case this could be achieved with flexible equalisation capabilities.

In terms of venue acoustics, it has been suggested that bottom-up practitioners tend to favour controlled acoustics in which reverberation (and so on) is minimised, while top-down practitioners are more inclined to regard unpredictable performance conditions as an interpretative challenge that, in a sense, justifies the very act of diffusion itself. These contrasting standpoints would, indeed, seem difficult to reconcile within the scope of a single sound diffusion system, but nonetheless, a compromise could be reached. Again, means to 'correcting' less-than-ideal listening environments could be implemented for the benefit of bottom-up practitioners; these might include phase correction, careful corrective filtering, calculated delays for individual loudspeaker feed signals, and so on.

Overall – as noted previously – the basis for evaluating and performing top-down works is essentially perceptual, whereas bottom-up works are conceived, realised, performed and evaluated according to conceptual or 'objective' criteria. This means that top-down and bottom-up practitioners are likely to expect different things from the sound diffusion system. For the top-down practitioner, the means of diffusion should perhaps offer a repertoire of real-time techniques suitable for shaping the musical material 'on the fly' in a perceptually effective manner. From this perspective, an intuitive and tactile interface allowing the diffuser to interact immediately and 'organically' with the musical material would seem appropriate. For the bottom-up practitioner, performance of the work is likely to be altogether more passive (perhaps taking the form of a transparent and objective presentation in which more attention is paid to the 'neutralising' of the performance context *prior* to diffusion than to the process of diffusion itself), or else precisely realised from a carefully preconceived 'score.' The latter case – like the bottom-up compositional process itself – demands a system whereby performance actions can be assessed according to their 'objective accuracy' (as opposed to their 'perceptual efficacy') and it is for this reason that bottom-up practitioners often express a preference for automated diffusion routines or procedures that can otherwise be consistently and accurately realised from a predetermined scheme. These two – in many respects diametrically opposed – methods collectively present additional challenges with respect to interface ergonomics, which will be discussed further in section 4.3.4.

It would be long-winded and impractical to individually describe all of the subtle nuances that differentiate top-down and bottom-up approaches to sound diffusion in detail. Reference to Table 4 (page 162) will give an indication of some of the broad philosophical and aesthetic contrasts and from these it should be possible to extrapolate more finite methodological differences. It is suggested that if top-down and bottom-up ideologies are to be fully supported by a sound diffusion system, then all of the factors described so far would have to be dynamically assignable during performance. This is to accommodate the fact that a concert programme could quite conceivably feature a top-down work immediately followed by a bottom-up work, each with quite distinct requirements from the diffusion system. Above all, this means that if a sound diffusion system is to be equally accommodating of both top-down and bottom-up methodologies, it must be enormously flexible, and able to facilitate (potentially fundamental) changes as necessary without disrupting the flow of the live concert programme.

In addition to these 'high level' considerations, there are additional criteria pertaining more to the logistics of sound diffusion itself. Sound diffusion is not simply a transparent means by which electroacoustic works are communicated to an audience, but an active practice in which hardware and software platforms are engaged, by a performer, for this purpose. In other words, a sound diffusion system is a creative framework in itself, and as such can be regarded, abstractly, as being similar to any other creative framework. As discussed in Chapters 1 and 2, implicit in every creative framework is a unique set characteristics, paradigms, and working techniques. The nature of the diffusion system itself (the processes involved in setting up and using the system; the control interface; the way in which hardware and software components interact; *et cetera*) will be a determining factor in evaluating the efficacy of that system as a platform for the communication of electroacoustic works. The following sections outline some of the criteria by which sound diffusion systems can be evaluated in these terms. Of course, the extent to which these criteria are met (or are not met) is very likely to impact upon the degree of success attainable in those 'higher level' criteria described previously.

4.3.3. Flexibility of Signal Routing and Mixing Architecture

Technically speaking, the aim of sound diffusion is to relay one or more coherent audio source sets – via an intermediate control interface – to one or more coherent loudspeaker sets. As described previously, there are many potential variables between works: the number and nature of audio sources; the number of channels contained within each CASS; the number and nature of the CLSs to be defined within the loudspeaker array; *et cetera*. In satisfying the two broad criteria given previously, it is therefore proposed that one of the single most advantageous features that a sound diffusion system can offer is flexible, dynamic signal routing and mixing capabilities.

What is ideally required is a signal routing and mixing architecture capable of handling an arbitrary number of audio input channels, to be mixed to an arbitrary number of output channels, via a user interface that allows the input and output channels, and input-to-output routings, to be conceptually grouped together (into CASSs and corresponding CLSs) in an arbitrary manner and ergonomically controlled (interface ergonomics will be discussed separately in section 4.3.4). Large numbers of input and (particularly) output channels are desirable, so long as this does not compromise flexibility. Such an architecture is represented diagrammatically in Figure 22, below.

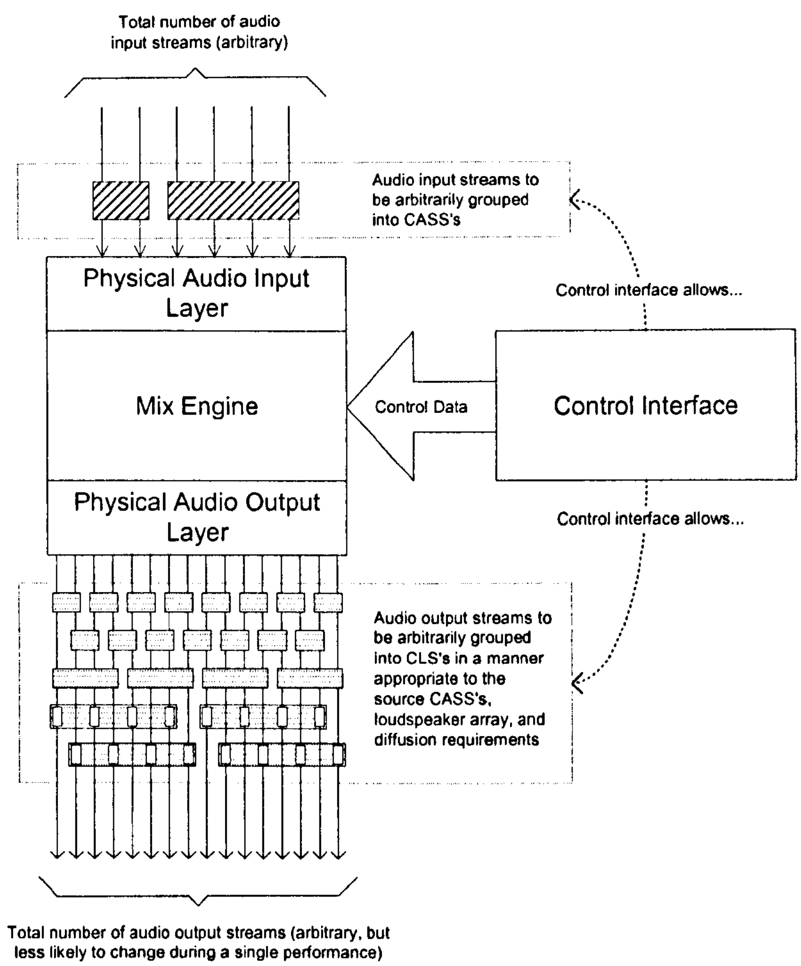

Figure 22. Diagrammatic representation of an ideal signal routing/mixing architecture for the purposes of live diffusion of electroacoustic music.

In this example, the total number of audio input streams (six) has been conceptually subdivided into two coherent audio source sets, one with two constituent channels (perhaps from a stereo pair of microphones) and one with four (perhaps from a quadraphonic fixed medium). As per section 3.8.3 (see page 140), this means that both two-channel and four-channel coherent loudspeaker sets are likely to be required within the loudspeaker array.[173] Of course the total number of input and output channels is an unknown variable and (particularly the former) may vary significantly from work to work within the context of a single concert programme.

This model is beneficial from the standpoint of criterion 4.3.1, because it does not make any assumptions with regard to the number of audio input and output channels, nor the ways in which these might be respectively grouped into CASSs and CLSs. This means that works demanding variable technical resources can easily be accommodated side-by-side. It also does not make any assumptions with regard to the relationship between the mixing architecture and the control interface, which – as will later become clear – is a criticism that can justifiably be directed at many existing sound diffusion systems. For the same reasons, this model is also beneficial with respect to certain aspects of criterion 4.3.2: CASSs and CLSs can dynamically be assigned in a manner appropriate for both uniphonic and pluriphonic diffusion methods (assuming, of course, that the loudspeaker array is able to accommodate all of the CLSs required), again with no restrictions imposed by the control interface itself.

4.3.4. Interface Ergonomics

'Interface ergonomics' relates to the physical logistics of interacting with a user interface (whether software or hardware, or a combination) as a means to achieving particular objectives in the process of sound diffusion. As simple examples, a top-down diffuser may wish to progressively 'widen' a stereo image by projecting it via consecutively wider-spaced loudspeaker pairs, or a bottom-up diffuser may wish to accurately articulate a carefully planned spatial orchestration. This

[173] The grouping of physical audio outputs into CLSs is also shown in Figure 22. Arbitrary groupings of the output channels into 'twos' and 'fours' (reflecting the nature of the input CASSs) are illustrated. The grouping lowest down on the diagram, for instance, shows that physical outputs 2, 4, 6, and 8, and 10, 12, 14, and 16, have been grouped together forming two four-channel CLSs. Of course, the actual groupings in any real scenario will depend on the shape of the loudspeaker array and the order in which physical audio outputs are routed to loudspeakers.

criterion evaluates the extent to which the user interface itself – by way of its particular architecture and paradigms – is able (or unable) to facilitate the realisation of such objectives.

If the control interface is to facilitate the complex demands outlined in the previous section, this raises some fairly serious questions with respect to how this can be achieved ergonomically, and in a manner that is suitable for live, real time, performance. Architecturally, matrix mixers (which will briefly be discussed in section 4.14) are able to fully support the ideal model presented in Figure 22, above, insofar as any input stream can be mixed, in any proportion, to any physical audio output. However, the number of physical controls (often rotary potentiometers) required for a matrix mixer with individual control over every input-to-output crosspoint can render the interface non-ergonomic for the purposes of live performance, that is, it would be logistically difficult to realise certain diffusion actions. This is, of course, indicates the necessity for a certain degree of compromise in the satisfaction of criteria 4.3.3 and 4.3.4

Additionally, one must consider the fact that diffusion equipment must, for the most part, be operated by performers in real time, and under reasonably stressful (live performance) circumstances. When this is the case, the control interface must ideally be fairly simple and intuitive, but nonetheless able to provide the performer with all of the necessary means to achieving the musical communication required *in real time* (except, of course, where preconceived automated diffusion is concerned).

4.3.5. Accessibility with respect to Operational Paradigms

This criterion is closely related to the previous, but is more concerned with the relative ease or difficulty experienced in the appropriation of interface paradigms by performers. To rephrase, this criterion evaluates the 'approachability' of the interface used to realise the diffusion of electroacoustic works, and the ease with which the necessary skills can be assimilated; the 'steepness' of the learning curve, in essence.

Implicit here is the notion of familiarity. If an interface is familiar to its user, then the range and quality of results attainable with that interface are likely to be fairly high. If the interface is unfamiliar, then it will be necessary for the performer to learn how to use it, and until this has happened the range and quality of results attainable will be limited. The difference between this criterion and the previous one rests in the fact that an interface can be highly ergonomic in the hands of a familiar user, but effectively useless in the hands of a less familiar user. Harrison makes the following observations:

> [The mixing desk fader is] an interface that practitioners of electroacoustic music tend to know fairly well. [...] Those of us who grew up making pieces with analogue tape *had* to use faders. You *had* to use the mixer because it was the only way of mixing stuff together, so you got to know what a fader felt like. So that's why it worked, historically, quite well, and has been the dominant one. [...] All these new, fancy things – gloves and motion sensors and so on – they all have their place, I think, but it's a question of how much time do we have to be able to get to control them sufficiently to *know* what we're going to do? I think that's the big issue. [...] The big problem with all of them is that every single possible configuration is a 'new instrument,' and an instrument takes time to learn how to play. At least with the good old fader we have a reasonable history of having learned to play it a bit. [...] But in a sense what you're looking at with everything you're doing is a range of values from *a* to *b*. A fader is actually a pretty efficient way of doing that, if you think about it.[174]

As implied by Harrison, the importance of the present criterion is reinforced by the fact that rehearsal time for concerts of live electroacoustic music is often minimal and, consequently, performers typically have very little time available to learn new paradigms. If an interface is unfamiliar, therefore, it is – on the one hand – desirable that performers should be able to assimilate the necessary skills quickly and easily. On the other hand, if an interface is quick and easy to master, the range of results attainable by more proficient users may be limited as a direct consequence, and this is – paradoxically – a case against interfaces that are too 'easy.' Kirk and Hunt stress the importance of

[174] Harrison (2004). "An Interview with Professor Jonty Harrison". Personal interview with the author, 8th September 2004. See Appendix 1 for full transcription.

mediating between these two opposing demands in the design of creative user interfaces; although they refer specifically to computer software interfaces, their point remains very valid in the context of sound diffusion:

> Most people would expect an ideal computer system [read: diffusion system] to be easy to use and easy to learn. However, nobody expects learning to drive a car to be easy. No musician would ever say that a violin or a clarinet was easy to use. Yet the subtlety of control and real-time interaction which can be shown by car drivers and instrumental performers is often astounding. By making the assumption that interfaces should be *easy* we are in danger of undervaluing the adaptation capabilities of human operators and thus limiting the potential human-computer interaction to a lowest common denominator of 'easy to use' commands. [...] A musician who has to work hard to produce a good sound is likely to be rewarded with appreciation from an audience. In contrast 'easy-play' instruments tend to produce 'cheap-and-easy' musical results.[175]

4.3.6. Accessibility with respect to Technological System Prerequisites and Support thereof

The 'technological prerequisites' of a sound diffusion system refers, basically, to the pieces of software and hardware required to build that system. If a system is accessible in terms of this, then it should be relatively straightforward for an institution to build the system in question. If a system is inaccessible in terms of its technological prerequisites, on the other hand, then the building of that system is likely to be less feasible or, in the worst case, impossible.

A number of discrete factors contribute to the relative accessibility or inaccessibility of diffusion systems in this context. Finance is a factor: if a diffusion system comprises expensive software and/or hardware components then it will obviously be inaccessible to any individual or institution without the financial means to obtaining these. Accessibility in terms of simple *availability* is also an important consideration. If a sound diffusion system consists of components that are difficult or impossible to obtain – perhaps they are 'custom' components or items that are not commercially or otherwise easily available – then that system will consequently be less accessible, regardless of any financial

---

[175] Kirk and Hunt (1999). *Digital Sound Processing for Music and Multimedia* (Oxford; Focal Press): 311-312.

considerations. The third sense in which a diffusion system can be evaluated in terms of this criterion concerns the level of technical support (access to documentation, shared knowledge, *et cetera*) available for the software and hardware components and this is, therefore, fairly closely related to the previous criterion. If system components are poorly supported in this respect, then it will be more difficult for performers to assimilate the necessary operational skills and the system will consequently be less accessible. Again, this is particularly true in the case of unconventional and/or custom-built components.

The importance of system accessibility in this respect has already been alluded to in the previous section. Kirk and Hunt (and, less directly, Harrison) liken creative interaction with a system to the practice of playing a musical instrument, noting that to become an accomplished instrumental player often involves years of disciplined training. This comparison is particularly valid in the context of sound diffusion, which is indeed a form of musical performance. In terms of the various factors described in the present section, traditional musical instruments (oboes, pianos, guitars, violins, drum-kits, *et cetera*) are reasonably highly accessible: financially, most teaching institutions (and, in some cases, individuals) are able to afford them; certainly, these instruments are readily available, and culturally well-supported in terms of tuition, relevant literature, performing ensembles, and so on. So, while these instruments may not be enormously accessible in the sense of item 4.3.5 – as observed by Kirk and Hunt, it cannot justifiably be argued that they are 'easy to play' – this is offset by a high degree of accessibility in the present sense.[176]

In any performance-related musical discipline, access to the instrument (in this case, the diffusion system) is important primarily because it

[176] In this respect, items 4.3.5 and 4.3.6 can be regarded as 'inversely proportional' to each other, insofar as a relatively low level of accessibility in terms of the one can be compensated for, to a certain extent, by a higher level of accessibility in terms of the other. A relatively high degree of accessibility in both senses would, of course, be ideal.

allows performers to practise in their own time. It would be completely unreasonable to expect a concert pianist to rehearse only on the day of the performance itself, yet this demand is – in the author's experience – fairly usual in the context of sound diffusion. This could be attributable to the fact that sound diffusion systems – as performance instruments – are relatively inaccessible in the present sense, and therefore often become available to performers only on the day of the concert. If systems were more accessible then institutions – perhaps even individuals – could build systems to rehearse with outside of the performance context, thus greatly increasing the potential for accomplished performance.

4.3.7. Compactness / Portability and Integration of Components

Performance venues in which equipment for the live diffusion of electroacoustic music is permanently installed are comparatively rare (although a few permanent systems will be described later in the chapter). It is therefore reasonable to state that the process of staging a performance of electroacoustic music will often entail the transportation of all of the necessary diffusion equipment to the venue, and setting up on location. If this is the case, it would seem advantageous for the system to be as portable as possible, and relatively quick and simple to set up. In addition, it is desirable that all of the equipment should be sufficiently robust to withstand the stresses sustained in transit, and in repeated setting-up and de-rigging throughout its life-cycle.

'Integration' implies the extent to which individual system components are 'normalised,' where appropriate, in terms of their relationships with one and other. If a system comprises components whose relationship within the context of the system as a whole is *fixed*, it would seem logical to integrate these. As a simple example, a CD player – rather than being transported as a separate unit – could be hard-wired into the physical audio input layer, assuming this is appropriate for the system in question. This would reduce the amount of time taken to set the system up and improve system reliability (to be discussed in section 4.3.8) by minimising the potential for signal routing errors, faulty cables, and so on. Care must be taken, however, to ensure that the process of integration does not jeopardise the flexibility of the system in terms of other criteria.

Aside from the obvious advantages, a system that is compact, portable, and well integrated might yield advantages with respect to some of the other evaluation criteria. If performers are to be able to practise in their own time, then a compact system that can be set up and de-rigged quickly and easily would be highly advantageous, for example.

The present criterion does not apply so obviously to diffusion systems that are designed specifically for permanent installation in a performance venue. It must be considered, however, that permanently installed diffusion systems – by definition – can only be used inside the venue within which they are installed. Unless a similar system can be built elsewhere (see criterion 4.3.6), this means that only performers with immediate access to the venue in question will have access to the system for practice. It seems logical to propose, therefore, that for a permanently installed diffusion system that is unavailable elsewhere to be of widespread use, it would have at the very least to 'score highly' in criterion 4.3.5 (i.e. it would have to be fairly quick and easy to learn).

4.3.8. Reliability

It should always be borne in mind that sound diffusion is a form of live performance, and this, of course, means that a diffusion system must take into account all of the issues inherent in live performance and real-time operation. An important consideration, in this respect, is reliability: if the technology is not stable in performance then it is not appropriate for the task in hand. Robustly-built hardware systems would seem the most likely candidates to score highly in this criterion, with systems incorporating software components (which can, in some cases, crash unexpectedly) potentially being less reliable.

It almost goes without saying, therefore, that software written specifically for the purposes of sound diffusion should be optimised for complete stability and reliability in performance. Assuming the software itself is well-written, unpredictable problems can arise when host computers of varying specifications are used; this can be presumed usual in the case of multi-purpose Windows-, Macintosh-, and Linux-based PCs, which are very rarely identical in software and hardware specification. If the undoubted advantages of software technologies are to be implemented in diffusion systems, therefore, the ideal in terms of reliability would consist in rigorously tested software, running on a system of known software and hardware profiles, and preferably dedicated *solely* to the purpose of running the diffusion software.

Another consideration in evaluating the reliability of diffusion systems (whether hardware, software, or a combination thereof) is the extent to which any malfunctions that do occur can be rectified quickly and easily within a performance context.

4.4. Criteria for the Evaluation of Sound Diffusion Systems: Summary

In summary, it is proposed that the following 'high level' criteria can be used to assess the effectiveness of sound diffusion systems as a means to publicly performing works from across the spectrum of electroacoustic music:

- 4.3.1: Ability to Cater for the Technical Demands of Live Electroacoustic Music
- 4.3.2: Ability to Cater for both Top-Down and Bottom-Up Approaches to Sound Diffusion

In the broadest possible terms, these two criteria have been designed for the purpose of evaluating the ability of sound diffusion systems to support electroacoustic works on a technological and aesthetic basis, respectively. Subordinate to these 'high level' criteria are several further criteria:

- 4.3.3: Flexibility of Signal Routing and Mixing Architecture
- 4.3.4: Interface Ergonomics
- 4.3.5: Accessibility with respect to Operational Paradigms
- 4.3.6: Accessibility with respect to Technological System Prerequisites and Support thereof
- 4.3.7: Compactness / Portability and Integration of Components
- 4.3.8: Reliability

It should already be reasonably clear that most of these criteria are strongly interrelated, and that the relationship between them is potentially complex. It seems unlikely that any diffusion system could perfectly satisfy all of the criteria simultaneously, because they represent the conditional factors of a multidimensional compromise that has no ideal solution. In other words, in satisfying the criteria, certain trade-offs will be more or less inevitable.

4.5. Description and Evaluation of Existing Sound Diffusion Systems

The following sections will seek – on the basis of information available in the literature and the author's own experience – to evaluate certain existing sound diffusion systems with respect to as many of the criteria given in section 4.3 as possible. Evaluation of existing diffusion systems cannot be absolutely systematic, as such. This would be enormously long-winded and, in any case, the literary information available is not always sufficient to permit such an approach. Nonetheless, it is proposed that the criteria should be of use to anyone wishing to undertake such a task, and with sufficient information and experience at their disposal to do so.

The purpose of this appraisal is not by any means to discredit existing diffusion systems in comparison with some unattainable ideal, but rather to broadly evaluate existing systems in terms of their appropriateness for the task of performing live electroacoustic music. Where limitations unduly interfere with this overall objective, possible improvements in the design of future systems will later be proposed. Equally, where advantageous features are identified, these will be acknowledged as features that could or should be incorporated in the design of future systems.

4.6. Generic Sound Diffusion Systems

As evidenced in the previous sections, the practice of sound diffusion presents a fairly unique, not to mention challenging, set of demands. Nonetheless, systems developed *specifically* for this purpose are surprisingly few, many performances of electroacoustic music (the majority, it is reasonably safe to suggest) being realised with what will be referred to as 'generic' sound diffusion systems. A generic diffusion system (GDS) comprises what might be considered 'every day' studio hardware. That is to say, it consists of an '*ad hoc* collection of equipment'[177] whose functionality is not uniquely devoted to live concert sound diffusion and was perhaps not originally designed or obtained for that specific purpose. Typically these systems are built around a conventional studio mixing desk that serves as the physical audio input layer, mix engine, control interface, and physical audio output layer. Such systems are very common, and are in use at the University of Glasgow[178] and at the Spatial Information Architecture Laboratory (SIAL) of RMIT (Royal Melbourne Institute of Technology) University,[179] to give but two examples. Two variations on this design are in common use.

4.6.1. Generic Diffusion System 1 (GDS 1): Description

A typical small system of this nature might consist of the following main components:

- Audio Source – most often a compact disc player or other stereo player.
- Studio Mixing Desk.
- Loudspeaker Array.

[177] Roads, Kuchera-Morin and Pope (2001). "The Creatophone Sound Spatialisation Project". Website available at: http://www.ccmrc.ucsb.edu/wp/SpatialSnd.2.pdf.
[178] University of Glasgow (2005). "Department of Music: Diffusion System". Website available at: http://www.gla.ac.uk/departments/music/studio/diffusion.html.
[179] SIAL (2002). "Sound Diffusion System". Website available at: http://www.sial.rmit.edu.au/Resources/Sound_Diffusion_System.php.

Figure 23 and Figure 24, below, illustrate the typical hardware configuration and routing capabilities, respectively, of such a system. Audio source channels are routed to mixing desk input channels, which are typically routed internally to the mixing desk's group buses. In a mixing desk with eight group channels – the commonly used Soundcraft Spirit Studio desk for example[180] – the 'left' audio source channel (such systems are inherently predisposed to stereophonic sound diffusion) would be routed to input channel 1 and this, in turn would be routed to group buses 1, 3, 5, and 7. The right audio source channel would be connected to input channel 2 and routed to groups 2, 4, 6, and 8. Thus, four copies of the stereophonic CASS are obtained. Group bus outputs are then routed appropriately to active loudspeakers – in this case the odd group channels are routed to 'left' loudspeakers, and the even channels to 'rights' – and the audio source can be diffused among stereophonic coherent loudspeaker sets using the group faders. This technique is described by Harrison as follows:

> An easy way to derive the multiple outputs needed is to use the 'groups' on a standard mixer – the stereo input from the DAT or CD is routed to all available groups and the relative levels achieved by 'playing' the group faders.[181]

As indicated in Figure 24, the group buses are in fact usually arranged into four stereo pairs, with the left-right balance for each channel sent to a stereo group bus being determined by the position of the pan pot for that channel.

[180] This mixing desk is used for sound diffusion at the University of Wales, Bangor and was, until recently, also used at the University of Sheffield Sound Studios. A slightly larger version was used in Performance Space 1 at the Sonic Arts Network *Expo 966* conference held in Scarborough, 17th – 20th June 2005.

[181] Harrison (1998). "Sound, Space, Sculpture - Some Thoughts on the 'What,' 'How' and 'Why' of Sound Diffusion". *Organised Sound*, 3(2): 123.

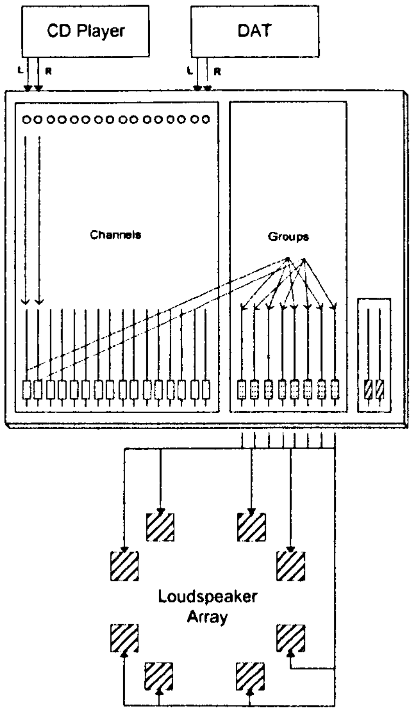

Figure 23. Set up and basic signal routing diagram for GDS 1.

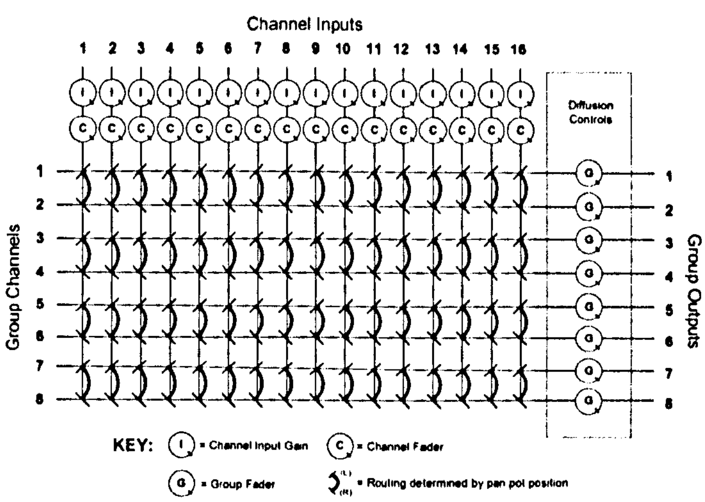

Figure 24. Schematic diagram illustrating the routing capabilities of a typical GDS 1 based on a 16/8/2 bus studio mixing desk.

4.6.2. Generic Diffusion System 2 (GDS 2): Description

A widespread variation of the system described above incorporates a 'splitter box' into the system between the audio source and the mixing desk. This system is illustrated diagrammatically in Figure 25 and Figure 26. The splitter box receives the audio source signals and duplicates them several times, relaying copies to its multiple outputs. These output streams are then routed directly to the input channels of the mixing desk, with loudspeaker feeds being taken directly from individual channel outputs. In this case the task of CASS duplication – particularly important in the case of pluriphonic diffusion – takes place outside the mixing desk. Diffusion is executed using the channel faders to control individual loudspeaker levels.

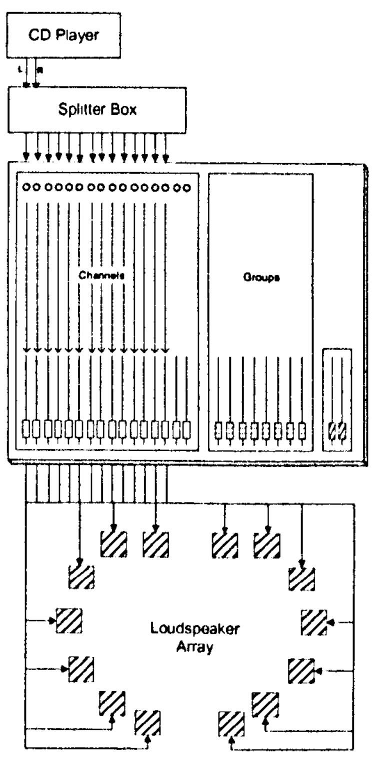

Figure 25. Set up and basic signal routing diagram for GDS 2.

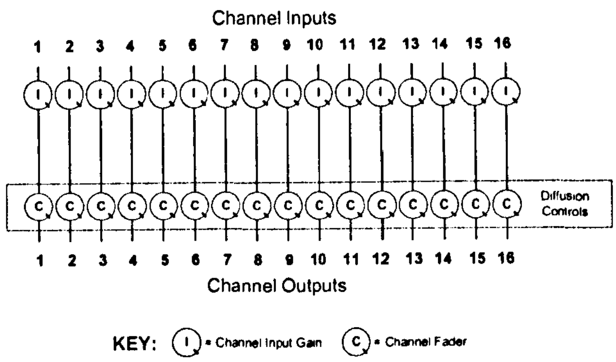

Figure 26. Schematic diagram illustrating the routing architecture of GDS 2.

4.6.3. Generic Diffusion System 1 (GDS 1): Evaluation

With respect to criterion 4.3.3, an advantage of this system is that the audio sources routed to the group faders can be switched. In performance situations this makes it relatively easy to deal with consecutive compositions whose audio streams originate from different sources: a CD player can be routed to input channels 1 and 2, a DAT player to channels 3 and 4, a stereo microphone pair to channels 5 and 6, an ADAT player to channels 7 to 14, and so on, and performers can alternate between audio sources reasonably quickly and easily, either by bringing up the levels of the appropriate channel faders or by switching between group bus routings, in both cases without the need for physical re-patching. This represents a reasonable degree of flexibility in terms of signal routing.

In terms of its mixing architecture (a further dimension of criterion 4.3.3), however, GDS 1 is less flexible. Although audio streams from multiple sources can be presented simultaneously (works for live instruments and tape, for example, could be accommodated to a certain degree), the extent to which these sources can be *independently diffused* is minimal. This is a function of the fact that the signal routing/mixing architecture and the control interface are intrinsically linked: the faders with which the performer controls the diffusion of audio sources to loudspeakers are also an integral part of the system that deals with the routing and mixing of audio source streams to physical outputs. Because there is only one fader per physical output, the only way that (say) two CASSs could be independently diffused would be via two separate sets of loudspeakers and corresponding faders. As well as being fairly unergonomic, this would also diminish the number of loudspeakers available for the diffusion of each CASS. This is a fairly serious problem given that the number of group channels (used as loudspeaker feeds in this system) available is likely to be fairly small to begin with. Indeed, a general disadvantage of GDS 1 is that the number of group channels available on the mixing desk limits the maximum number of loudspeakers in the array: for example, eight group channels allows for a total of only eight loudspeakers.[182] While this allows reasonable scope for the pluriphonic diffusion of stereo works,[183] CASSs with more channels would be subject to increasing restrictions, with CASSs of five or more channels effectively being limited to uniphonic presentation.

The fact that the amplitude levels of fixed input-to-loudspeaker routings are each controlled individually with a single fader (in some cases simultaneous control of a stereophonic pair with one fader may be possible) also presents problems in terms of interface ergonomics (criterion 4.3.4). In combination with the fact that faders are almost invariably arranged in a horizontal row, this means that certain diffusion actions will be physically difficult to execute. This problem would appear to be fairly widely recognised:

> The speed at which decisions can be implemented and the dexterity of control required when operating equipment clearly influences performance practice. Given the practicalities of very little rehearsal time in often

[182] In practice, an additional pair of loudspeakers can be obtained by using the (normally stereo) master outputs. Further channels can also be utilised via a somewhat clumsy method of 'copying', by connecting channel sends to the inputs of unused channels and using the direct outputs of these channels (assuming, of course, that these are provided) to drive loudspeakers. Additionally, auxiliary buses can be used as independent channels, although this is obviously far from ideal, as the levels thereof are normally controlled with rotary potentiometers as opposed to the more familiar (and perhaps more appropriate) sliders. Clearly, none of these techniques are particularly convenient.
[183] Although, recalling the quotation given in section 3.12.1 (page 151), Harrison regards eight loudspeakers as the *minimum* for the diffusion of stereo works.

> inappropriate venues, this basic setup can be either extremely limiting or (in the case of a large system) highly intimidating...[184]
>
> **JM**: In your experience, how much is your diffusion of any given piece restricted – or indeed enhanced – by the fact that you're having to do it on a 'one fader, one loudspeaker' basis, where the faders are arranged [from left to right]?
>
> **JH**: Well, inevitably it is influenced by that, because there are physical constraints. You can of course re-patch or re-jig and get it to be better, and easier to do certain things. I've done a lot of things in France, for example, where they tend to start at one end of the mixer with the speakers that are the furthest away from you in front, and then [gestures from left to right] the next pair, the next pair, the next pair, and this end [gestures to the right] is at the back. That's absolutely tickety-boo if you want to do lots of front-to-back or back-to-front things, but if you want to get the most significant speakers near the audience to be going [makes antiphonal sounds and hand gestures], you'll find that there's a pair there, a pair there, a pair there, and a pair there [indicates four completely separate areas of an imaginary mixing desk]; you can't get that interaction, which is why we don't do that in BEAST. We put those main speakers in a group on the faders so that you can do it easily. But then of course it means that if you want to go front-to-back, it's awkward. So, nothing's perfect! Unless you could suddenly re-assign the faders mid-diffusion. *You* [referring to the M2 Diffusion System, to be described in Chapter 5] could do that, couldn't you? [Laughs].
>
> **JM**: The bottle-neck as I see it, is the fact that most systems work on a 'one fader, one loudspeaker' basis. While you'll have the advantage of knowing exactly what everything does (that's the positive side), the negative side is that you essentially have an arbitrarily organised set of controls that are fixed. So there will be things that are just going to be very difficult to do.
>
> **JH**: That is true.[185]

With this architecture and control interface, issues of ergonomics are exponentially compounded in the case of works for more than two channels.

> If you've got an eight-channel source, and every channel of the eight has a fader, how do you do crossfades? You haven't got enough hands![186]

This indicates that, as far as pluriphonic diffusion goes, GDS 1 is, in practical terms, easier to use, and more flexible, for single-CASS works with a small number of channels: given current conventions, this basically equates to stereophonic fixed medium works. Works

---

[184] Moore, Moore and Mooney (2004). "M2 Diffusion - The Live Diffusion of Sound in Space". *Proceedings of the International Computer Music Conference (ICMC) 2004.*
[185] Harrison (2004). "An Interview with Professor Jonty Harrison". Personal interview with the author, 8th September 2004. See Appendix 1 for full transcription.
[186] Ibid.

comprising a single CASS of more than two channels are, essentially, limited to uniphonic presentation, as are works incorporating more than one CASS.

As noted, many of the problems described can be directly attributed to the fact that control of the amplitudes of each loudspeaker is on a one-fader-to-one-loudspeaker basis, and therefore actions involving multiple loudspeakers will necessarily involve potentially complex (or, at worst, impossible) interactions with multiple faders. Nonetheless, this does afford the diffuser a relatively large degree of 'low level' control and therefore (relating to section 4.3.5) offers the potential for accomplished and virtuosic performance (particularly in the case of stereophonic works), notwithstanding certain issues relating to interface ergonomics.[187] The potential for virtuosity is augmented by the fact that the mixing desk fader – as observed by Harrison in section 4.3.5 – is an interface familiar to most practitioners of electroacoustic music (and also in other areas) and is therefore accessible as an operational paradigm.

There are a number of further issues arising from the appropriation of conventional mixing desks (which were not originally intended for this purpose) in GDS 1. The effective presence of two input gain controls per channel (namely the input gain itself, and the channel fader – this is best observed in Figure 24) is clumsy, and can lead to unpredictable signal levels in concert, with performers potentially adjusting the overall input gain with two separate controls. Additionally, because audio mixing desks almost invariably provide more channel strips than group buses, so diffusion systems utilising mixing desks in this way can exhibit an over-abundance of audio source inputs, as well as an inadequate number of group bus outputs for use as loudspeaker feeds.

[187] The keys of a piano, after all, are also arranged horizontally in a row, meaning, of course, that certain combinations of notes will be physically impossible to play, and comparable 'limitations' can be observed in all creative frameworks. It is rarely suggested (in the author's experience) that these limitations render the instrument ineffective; on the contrary, the potential for subtle and virtuosic performance on such instruments is widely acknowledged, perhaps in part *because* of the difficulty associated with playing them.

Overall, GDS 1 benefits from a high degree of accessibility in terms of both operational paradigms and technological system prerequisites (this will be discussed more fully in section 4.6.5). A reasonable degree of flexibility with regard to signal routing is offered, but this is let down by an overly restrictive mixing architecture, which is too closely integrated with the control interface. Consequently, GDS 1 presents fairly serious problems with regard to interface ergonomics. The various technical demands of electroacoustic music are reasonably well-supported (criterion 4.3.1) insofar as multiple technologies can be accommodated and the signal routing architecture allows for these to be dynamically assigned. They are less well supported in respect of multiple and/or variable CASS demands, the system only really offering any degree of flexibility in the pluriphonic diffusion of single-CASS stereophonic works. Accordingly – with respect to criterion 4.3.2 – it could perhaps be argued that GDS 1 exhibits a slight bias towards top-down works and diffusion methods. On the other hand, the system presents no major difficulties for the *uniphonic* presentation of works, which, in any case, tends to be favoured by bottom-up practitioners. It can therefore be proposed that GDS 1 does not exhibit any substantial top-down or bottom-up bias, but – at the same time – neither does it offer an enormous degree of flexibility or creative scope with respect to either of these contrasting performance practice methodologies.

4.6.4. Generic Diffusion System 2 (GDS 2): Evaluation

GDS 2 has the considerable advantage of being able to utilise all of the channels on the mixing desk as speaker feeds (provided the splitter box has a sufficient number of output channels, that is) and can therefore potentially drive a substantially larger loudspeaker array than GDS 1. This is simply due to the fact that normal studio mixing desks almost invariably provide more input channel strips than group buses. The system also has the advantage of offering independent EQ on every effective output channel, which might be useful in compensating for loudspeakers with non-linear frequency responses, or indeed for 'colouring' loudspeakers with uniform frequency response, should this be desired. This adds a dimension of flexibility with respect to both top-down and bottom-up methodologies that cannot be obtained with GDS 1 (which only offers EQ on the audio source inputs assuming a fairly standard mixing desk is being used). Additionally, this configuration has only one input gain control per channel (see Figure 26, comparing this with the two effective input gain controls superfluously featured in GDS 1, illustrated Figure 24), thus reducing the possibility of unpredictable signal levels during performance and improving overall system reliability.

In these respects a generic diffusion system incorporating a splitter box can potentially make more efficient use of the resources available than one without. However, GDS 2 does still exhibit a certain degree of hardware redundancy insofar as the group and master sections of the mixing desk will normally not be utilised. Consequently, the pan-pots on each channel strip are also effectively useless. Additionally, as Harrison observes, this setup is only possible 'if the mixer has direct outputs on all channels, [and] lower-cost desks tend not to have this feature.'[188]

Although it yields a number of advantages, GDS 2 also presents some considerable disadvantages, most notably in that (mainly owing to its reliance on a splitter box) the signal routing and mixing is far less flexible. In addition to the difficulties arising from a control interface that is integral to the signal routing/mixing architecture (these were described in the previous section and GDS 2 presents exactly the same problems), in GDS 2 the input-to-output routings must effectively be *hard-wired* when the system is set up (in GDS 1 the input-to-output routings are at least switchable via the group bus architecture). Each channel fader is part of a simple series circuit comprising audio input, in-line attenuator (fader), and audio output (again, see Figure 26), and

[188] Harrison (1998). "Sound, Space, Sculpture - Some Thoughts on the 'What,' 'How' and 'Why' of Sound Diffusion". *Organised Sound*, 3(2): 123.

owing to this direct and physical routing of inputs to outputs (there is no abstraction whatsoever), the use of multiple audio sources in such a system is highly problematic. In practical terms (unless a particularly sophisticated splitter box with switchable input-to-output routings is being used) the system is limited to a single audio source if physical re-patching during performance is to be avoided.

Ergonomically, the control interface offered in GDS 2 is practically identical to that of GDS 1, because both systems utilise normal studio mixing desks. Although GDS 2 can potentially host a loudspeaker array large enough for the pluriphonic diffusion of multichannel works (its architecture allows for this; the same cannot realistically be said of GDS 1), the fact that each source-channel-to-loudspeaker amplitude must be controlled with an individual fader makes this logistically extremely difficult, as noted in the previous section. Stereophonic works, however, can be pluriphonically diffused reasonably easily and – like GDS 1 – the system presents no real problems for the uniphonic presentation of multichannel works, except that this is likely to necessitate physical re-patching unless all of the programme items utilise the same audio source technology. What this implies is that the benefits of GDS 2 – mainly, the fact that a larger loudspeaker array can be utilised – are only really advantageous in the pluriphonic diffusion of stereophonic works, meaning that GDS 2 is fairly biased in favour of top-down diffusion techniques.

4.6.5. Generic Diffusion Systems in General: Evaluation

An important advantage of both of the generic sound diffusion systems is that they utilise commonplace studio hardware, and therefore should be easily and conveniently within the capabilities of even fairly modestly equipped studios with comparatively little additional outlay. Generic systems are therefore accessible as per criterion 4.3.6, and indeed it seems very likely that this is one of the main reasons for their existence and continued use. These systems are fairly 'modular' in nature, insofar as additional components can be added with relative ease. This means that various different technologies can be accommodated (in accordance with criterion 4.3.1) but this is not so much a result of innovative design as a convenient characteristic of communication via the standardised protocols of encoded audio. Such systems are often, therefore, neither particularly portable nor integrated (criterion 4.3.7), with individual components having to be transported separately and time-consumingly assembled on location. This often necessitates an enormous amount of cabling (particularly for GDS 2, where a cable will be required to connect each splitter-box output to a channel input: this represents an obvious aspect in which optimisations could be made) as well as increasing the possibility of routing errors, forgotten components, and so on.

In both GDS 1 and GDS 2 there is the potential for a relatively high level of redundancy with respect to the hardware resources used. The root of this problem can be traced to the fact that mixing desks are designed and built with a view to mixing multiple audio streams into (usually) a two-channel stereo mix, whereas with sound diffusion, the aim is (in most cases) to take a finite number of audio channels and project these into a *greater*, or at least equal, number of audio channels. In this respect, the use of studio mixing desks for sound diffusion can be viewed as something of a misappropriation. As noted by Harrison:

> The drawback with [...] [GDS 1] is that even 8-group desks normally have 16 or more input channels, of which only 2 are really needed for the stereo input. If the mixer has direct outputs on all channels [...], then input channel faders can be used for diffusion, but this necessitates splitting the stereo signal out from the source (via parallel boxes, for example) and running several left/right signal cable pairs into successive pairs of input channels [GDS 2]. BEAST has developed an elegant way of achieving what is effectively a 'mixer in reverse' (2 in/many out) by using a switching matrix through which any incoming stereo signal can be routed to any pair of outputs. [This will be described later, in section 4.8].[189]

In addition, mixing desks can be heavy and cumbersome to transport – further diminishing portability – and this seems somewhat inefficient

[189] Ibid.

given that many of the hardware features will not, in fact, be used during performance.

Although the mixing desk fader is highly accessible as an operational paradigm (this is a definite advantage), the paradigm itself is not particularly ergonomic, owing largely to the fact that the control interface is also an integral part of the signal routing and mixing architecture, which (as already observed) is not designed specifically for the purpose of sound diffusion and accordingly falls short of meeting certain demands. Works comprising multiple CASSs and/or multichannel CASSs both present particular difficulties in this respect.

Overall, the single greatest advantage of generic diffusion systems rests in the very fact that they are generic. Put simply, most institutions will own all of the necessary hardware anyway, and this can be assembled in a manner that is (more or less) appropriate for the concert programme to be staged. Ultimately, however, the technology has not been designed specifically for sound diffusion, and is therefore not as well suited for this purpose as might be desirable.

4.7. The Acousmonium

4.7.1. Description

The Acousmonium was devised by François Bayle and Jean-Claude Lallemand at the Groupe de Recherches Musicales (GRM), hosting its first public performance in 1973.[190] It is this system that first gave rise the expression 'orchestra of loudspeakers,' using up to (and sometimes over) eighty loudspeakers in any single performance. Relatively little published literature relating directly and specifically to the Acousmonium exists in English. The following description is, therefore, based mainly on secondary sources and personal communications with the author.

[190] The first piece performed was François Bayle's *L'Experience Acoustique*: GRM (2004). "INA-GRM: Key Dates". Website available at: http://www.ina.fr/grm/presentation/dates.en.html.

One notable characteristic of the Acousmonium is its use of asymmetrical loudspeaker arrays; many other systems tend to favour strictly symmetrical arrangements. The Acousmonium also incorporates loudspeakers of different timbral characteristics as 'soloists,' an effect that is augmented by the physical arrangement of loudspeakers in a manner comparable to the instrumental sections in an orchestra. Harrison further observes that the loudspeakers are 'positioned in a mainly frontal array on the large stage of the Salle Olivier Messiaen,'[191] further reinforcing the orchestral analogy. Accordingly, it has been observed that the Acousmonium specifically draws attention to the loudspeakers as 'performers,' whereas many other systems, including the Cybernéphone, (which will be described in section 4.11) attempt to draw the audience's attention *away* from the loudspeakers during performance.[192]

Notwithstanding this (fairly subtle) distinction, others *liken* the Acousmonium to the Cybernéphone. According to Stevenson, both of these systems are similar insofar as:

> Intensity panning, reverberation, equalisation and loudspeaker design or 'registers' [are used] to modify the *perception* of individual sources and pre-mixed tracks. Loudspeakers are distributed throughout the auditorium. *The unique acoustic properties of these systems are exploited* in much the same way as a musical instrumental ensemble. It is [Stevenson's] impression that *adapting a piece for these ensembles* would be considered similar to the process of re-orchestration of instrumental music. [My italics.][193]

Stevenson's observations (particularly the italicised sections) appear to indicate that the Acousmonium and Cybernéphone are similar insofar as both seem to engender top-down approaches to the diffusion of sound. It will later be proposed that – certainly in the case of the Acousmonium, at least – this is indeed true.

[191] Harrison (1999). "Diffusion - Theories and Practices, with Particular Reference to the BEAST System". *EContact!* 2(4). Electronic journal available at: http://cec.concordia.ca/econtact/Diffusion/Beast.htm.
[192] NTMusic UK (2004). "Magnetic Tape Distribution". Website available at: http://www.ntmusic.org/proj_dmu12.htm.
[193] Stevenson (2002). "Spatialisation, Method and Madness - Learning from Commercial Systems". *Australasian Computer Music Conference* (Melbourne; 6th - 8th July 2002).

Diffusion is executed using a mixing-desk-style interface with a fader-based control paradigm. With up to eighty loudspeakers, control of the diffusion on a one-fader-to-one-loudspeaker basis would clearly be difficult. Therefore, according to Garro, the Acousmonium allows input signals to be routed to loudspeakers on a 'one-to-many' basis.[194] In this way, one fader can be used to control the amplitude of an input channel routed to multiple loudspeakers. It is safe to assume that, in the system's original 1973 incarnation, this signal routing architecture – if indeed it was available at this time – must have been facilitated in hardware; it is unclear as to whether or not subsequent software additions have been made.

4.7.2. Evaluation

Although there is no reason to believe that it is not technically capable of accommodating a reasonably wide variety of electroacoustic works, the Acousmonium – more so than either of the systems described previously – demonstrates a clear and considerable bias towards top-down methods, and this is due mainly to the nature of the loudspeaker array. The use of loudspeakers that 'colour' the sound is normally contrary to the bottom-up ethos of transparent reproduction of the electroacoustic work, and is also objectionable to some top-down practitioners. The asymmetrical arrangement of loudspeakers potentially makes it difficult to accommodate works that demand a specific or symmetrical loudspeaker array, for example multichannel fixed media. Because bottom-up works are more likely to make use of multichannel media, this represents another respect in which bottom-up methodologies are marginalised by this system. This is, of course, completely in-keeping with the top-down tradition, which – as described in Chapter 2 – is descended from *musique concrète* and has its roots in the GRM itself.

[194] Diego Garro (2005). Personal communication with the author.

Nonetheless, the ability to control the amplitude levels of multiple loudspeakers with a single fader is most certainly an advantage from a number of perspectives. Firstly, and most obviously, interface ergonomics are greatly improved. Secondly, support for multichannel works is improved: Garro reports that audio source channels can be individually assigned to different groups of loudspeakers,[195] meaning that multichannel fixed-medium works (which will normally demand a fairly specific CLS shape) can make selective use of the asymmetrical loudspeaker array. Given the large size of the loudspeaker array, this architecture might also permit simple pluriphonic diffusions of multichannel works, although in this case interface ergonomics may again become an issue. The system does have the advantage of being based around the familiar fader, and this compensates, to a certain extent, for the fact that it is not particularly portable.

Overall, the Acousmonium appears to be technically more flexible than either of the generic systems, but has a markedly distinct aesthetic bias. This bias is far more favourable to the top-down ethos than the bottom-up, but even within the scope of top-down practice the use of asymmetrical loudspeaker arrays with non-linear frequency response can be subject to strong criticism. Therefore, although apparently fairly flexible, the Acousmonium is only ideally suited to those practitioners who subscribe to the reasonably specific set of aesthetic values that it engenders.

4.8. The BEAST Diffusion System

4.8.1. Description

The diffusion system used by Birmingham Electro-Acoustic Sound Theatre (BEAST) includes a custom-built DACS 3D diffusion console, which Harrison describes as 'a mixer in reverse (two in, many out)'.[196]

[195] Diego Garro (2005). Personal communication with the author.
[196] Harrison (1998). "Sound, Space, Sculpture - Some Thoughts on the 'What,' 'How' and 'Why' of Sound Diffusion". *Organised Sound*, 3(2): 123.

A similar console is also in use at Keele University.[197] This console offers twelve channel inputs and thirty-two outputs, connected via a matrix of input-to-output routing buttons which can be switched on or off in stereo pairs. Additionally, the left-right orientation of each stereo routing can optionally be reversed, resulting in the routing schematic given in Figure 27, below. Each channel has an individual input gain control and insert point, and diffusion is realised using the thirty-two output channel faders. A more detailed specification for the 3D console is available via the DACS website.[198]

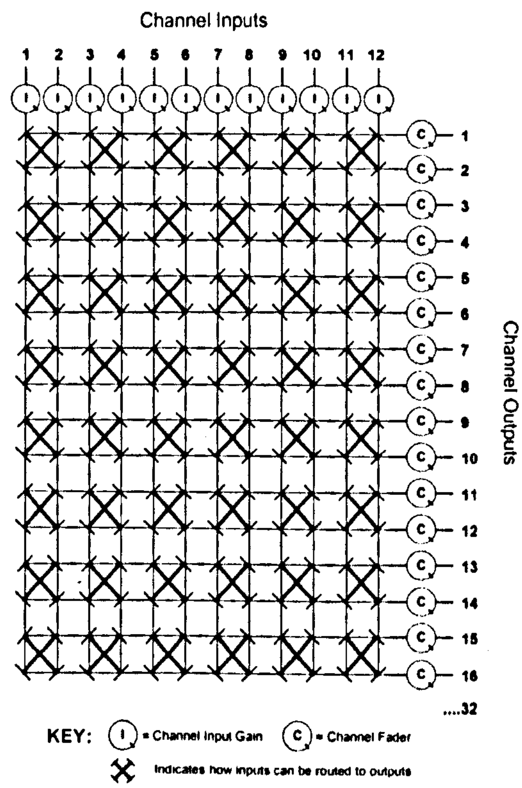

Figure 27. Simplified schematic diagram illustrating the signal routing capabilities of the DACS 3D Sound Diffusion Console

4.8.2. Evaluation

As confirmed by Harrison (see quotation given in section 4.6.5), the DACS 3D diffusion console was developed in response to practical

[197] Diego Garro (2005). Personal communication with the author.

[198] DACS (Unknown). "3D Sound Diffusion Console". Website available at: http://www.dacs-audio.com/Custom/3d_console.htm.

experience with the generic systems described in section 4.6, incorporating the beneficial features of each of the two variants and addressing some (but not all) of the problems. Input channels can be routed to output buses with the same degree of flexibility offered in GDS 1 (and lacking in GDS 2). Additionally, the console allows for diffusion across a large loudspeaker array, as in GDS 2, but without compromising the flexibility of the signal routing. Incorporating hardware designed specifically for the purpose of sound diffusion, the BEAST system represents an enormous improvement with respect to criterion 4.3.3 in comparison with the generic setups, effectively dispensing with the hardware redundancy previously cited as a flaw of the latter. Different audio sources can be selected reasonably quickly and easily with the input-to-output routing buttons, and the ample twenty-four audio inputs ensures that this is unlikely to necessitate physical re-patching except in fairly extreme circumstances.

However, the system presents precisely the same difficulties with respect to interface ergonomics that were observed in relation to the generic setups, insofar as the diffusion console operates on a strictly one-fader-to-one-loudspeaker basis. To reiterate, this can make the pluriphonic diffusion of multichannel works, in particular, logistically difficult. Indeed, the system would certainly appear to have been designed specifically with the diffusion of stereophonic fixed medium works in mind: routings are made in stereo pairs, and the typical BEAST loudspeaker array consists of multiple loudspeakers distributed throughout the venue in stereo pairs in the manner illustrated in Figure 14 (page 141). Accordingly, the diffusion of works with different technical demands – a greater number of fixed medium channels or live microphone feeds for instance – can become problematic to a certain extent. Although the BEAST system offers a certain degree of input-to-output routing abstraction (the routings are not 'hard-wired' as they are in GDS 2), problems can still be attributed to the fact that the routing/mixing architecture and control interface are too closely interrelated. Nonetheless, it can always be argued that the mixing desk fader is an accessible paradigm with respect to criterion 4.3.5 and this is, of course, an extremely important consideration.

In terms of portability, the BEAST system is broadly similar to the generic systems. Although the diffusion console undoubtedly represents a far more elegant solution (particularly in comparison with the cumbersome splitter box incorporated in GDS 2), audio sources – for example – are still provided by external units that must be integrated into the system in the performance venue. However, the DACS 3D is ruggedly built, regular transportation between venues presumably having been an explicit design consideration, and this cannot always be said for conventional studio mixing desks. In this respect, the BEAST system is potentially more reliable than either of the generic systems.

The BEAST system also benefits, to a certain extent, from the fact that audio source technologies can be integrated into the system in a manner comparable with normal studio mixing hardware, thus making the system a broadly familiar – and therefore accessible – platform to work with. The only 'non-standard' component is the mixing console itself, which currently must be built to order by DACS Limited, Gateshead. At the time of writing this can be done for a total cost of £5915, the order taking around two months to complete.[199] Although this is not, perhaps, as readily accessible as standard commercial mixing hardware, there are a number of advantages. Firstly, the desk has been designed specifically for the purpose of sound diffusion and is therefore generally more appropriate for the task. Secondly, the exact specification – in terms of the number of inputs and outputs – can be tailored to suit the particular requirements of the individual or institution in question.

Overall, the BEAST diffusion system is far more appropriate than the generic systems in a number of respects, as a platform for the live presentation of electroacoustic music. The technical demands of the

[199] Douglas Doherty, Managing Director of DACS Limited (2005). Personal communication with the author.

idiom are fairly well accommodated, mainly by virtue of the signal routing architecture, which is simple and effective and therefore well-suited to live performance. Like GDS 1, the BEAST system does not exhibit any overly prohibitive bias towards top-down or bottom-up techniques (in comparison with, for example, the Acousmonium, which is enormously biased in favour of the top-down), although it can certainly be argued that the top-down ethos is favoured to a certain extent. Like the generic systems, however, the biggest criticism is relatively poor interface ergonomics, stemming from the fact that diffusion must be executed on a one-fader-to-one-loudspeaker basis. This means that the capabilities of the system can only really be fully exploited in the case of pluriphonic diffusions of stereophonic fixed medium works, hence the slight top-down bias. Although bottom-up techniques are by no means institutionally excluded (this argument could, tentatively, be made with reference to the Acousmonium), neither are there any features that are of particular and specific use to bottom-up practitioners.

4.9. BEAST's Octaphonic Diffusion Experiments

4.9.1. Description

It has been noted that the standard BEAST diffusion system – in common with the generic systems – presents certain difficulties in the pluriphonic diffusion of more-than-stereo works. In response to this, experimental research into the pluriphonic diffusion of octaphonic fixed medium works (that is, the presentation of an eight-channel CASS via multiple eight-channel CLSs) has been carried out by BEAST. At the time of writing, these experiments have not been documented elsewhere.

In an interview with the author,[200] Harrison described an array comprising eighty loudspeakers arranged in ten eight-channel CLSs

[200] Harrison (2004). "An Interview with Professor Jonty Harrison". Personal interview with the author, 8th September 2004. See Appendix 1 for full transcription.

throughout the hall, roughly as illustrated in Figure 28, below. Six of the CLSs – 'main,' 'close,' 'floor,' 'higher,' and 'gallery' and 'ceiling' – were arranged in a fairly standard circular/ovoid octaphonic shape, with the speakers inclined inwards towards the centre. In the 'stage' array, all eight loudspeakers pointed towards the audience and were slightly raked so that those further away were positioned higher than those closest to the audience. The 'central' array comprised eight small loudspeakers arranged in a circle and pointing outwards. The 'ferris' array formed a vertically-oriented circle of eight loudspeakers around the audience on a left-to-right axis, like a Ferris-wheel, with each of the speakers pointing inwards. A second vertical array – 'rose' – was set up on the far wall of the venue resembling a rose window in a church, with each of the speakers facing towards the audience.[201]

[201] As an aside it is interesting to note that, although each CLS is slightly different in formation, and some are more different than others, they all conform to the appropriate 'shape' for this kind of coherent loudspeaker set. That is, they all meet the criteria – in Harrison's estimation, which is oriented toward the top-down end of the spectrum – for the accurate reproduction of the coherent audio source set(s) in question. This clarifies and reinforces certain points that were made in sections 3.5 to 3.8.

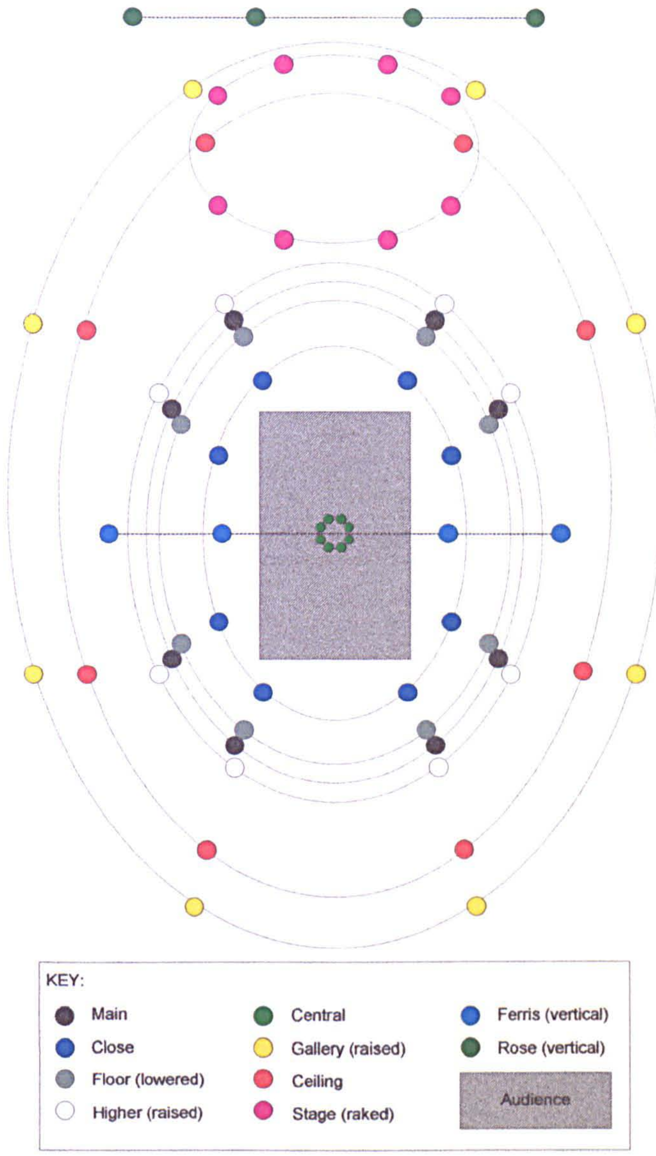

Figure 28. Approximate loudspeaker positions for the BEAST octaphonic sound diffusion experiments.

Such a system would obviously be completely unmanageable with a fader controlling the level of each individual loudspeaker. Accordingly, diffusion was controlled using a MIDI fader box in conjunction with an intermediate software interface programmed in Max/MSP, where each MIDI fader controlled the amplitude level of a group of eight loudspeakers. This, of course, supports the notion of coherent audio source sets and coherent loudspeaker sets as useful concepts in the context of sound diffusion. The result was a ten-fader interface controlling diffusion to an eighty-strong loudspeaker rig. In addition, an eleventh 'virtual' fader was created that inverted the octaphonic image in the 'main' array, that is, input 1 was routed to output 8, input 2 to output 7, input 3 to output 6, and so on. This demonstrates a degree of input-to-output routing flexibility that is extremely difficult to achieve with standard audio mixing hardware, with the exception perhaps of matrix mixers, which will be discussed later.

4.9.2. Evaluation

It must first of all be noted that the system described in the previous section is very much experimental and has not (at the time of writing, and as far as the author is aware) been used for any public performances or further experiments to date. The system was not presented to the author as a 'finished product,' and can therefore not be justifiably evaluated as such. Nonetheless, an evaluation in terms of the criteria given in section 4.3 – as far as this is reasonable – will prove useful.

The main advantage that this system offers is that the signal routing/mixing architecture and control interface are highly abstracted from each other. The same cannot truly be said of any of the systems described previously, where the control interface, at some level, is an integral part of the routing and mixing architecture itself. Here, the control interface 'merely' provides data relating to the position of each fader, and this data can ostensibly be interpreted in any way by the software mix engine. In this case, data relating to the position of each individual fader is used to simultaneously control the output levels of an eight-channel coherent loudspeaker set, with appropriate CASS-channel-to-loudspeaker routings also being facilitated in software. This is conceptually interesting as it foreshadows and supports the notions of CASS and CLS in sound diffusion, insofar as decisions must be made with regard to the grouping of audio source streams and corresponding loudspeakers. The benefits of this are particularly evident in the 'virtual' array in which input-to-output routings are dynamically inverted in software. This particular specification is, of course, exclusively suited to the pluriphonic diffusion of eight-channel works and can therefore be regarded, on the whole, as an extension of the top-down ethos broadly advocated by BEAST.

Clearly this setup presents vastly improved interface ergonomics: a cross-fade between two CLSs can be achieved with the manipulation of only two faders as opposed to the practically impossible sixteen that would be required to realise the same results with many other systems. This system also has the advantage of presenting the performer with a simple and (usually) familiar control interface – a bank of faders – whose operation is basically similar to that of a conventional studio mixing desk. Accordingly, the system meets criteria 4.3.4 and 4.3.5 reasonably well, mediating the potential 'trade-offs' inherent in these criteria in an effective manner.

It could also, of course, be argued that – in comparison with the potentials of abstracted software control – the system does relatively little, adhering quite closely to the somewhat limited model presented by conventional studio mixing hardware. However, it must be recalled that this is a basic experimental system designed only to explore fairly specific creative objectives. Indeed, this approach – using Max/MSP as a platform for the implementation of experimental software mix engines with abstracted physical control – could prove to be an exceedingly effective (not to mention relatively quick and straightforward) way of prototyping new sound diffusion paradigms, and much research could be undertaken in this way. Because the system incorporates fairly standard components (many institutions with an interest in electroacoustic music will have easy access to Max/MSP, MIDI fader modules, multichannel audio interfaces, *et cetera*) and is reasonably loosely specified, it is therefore accessible according to criterion 4.3.6, and such research can potentially be widely engaged with.

Accordingly, it can be said that the approach adopted by BEAST in these experiments is one that could be taken much further, with potentially enormous benefits in the field of sound diffusion. As it stands, the system is not, perhaps, as reliable as might be desirable in a live context: there is no indication, for instance, that the software has been extensively tested. Neither is the system as compact and integrated as it could be. Additionally, in this particular instance, there is a demonstrable bias in favour of the top-down agenda, to the practical exclusion of bottom-up techniques. However, such limitations are surely to be expected in an experimental research context, and the model adopted would certainly allow for these issues to be addressed in the future.

4.10. The Creatophone

4.10.1. Description

The Creatophone project is based at the Center for Research in Electronic Art Technology (CREATE) at the University of California, Santa Barbara, in the USA. Research in progress at CREATE expresses an interest in surpassing the 'one-fader-to-one-loudspeaker' paradigm with grouped loudspeakers, thus addressing issues similar to those tackled by the Acousmonium and by BEAST's experimental octaphonic system:

> The first full realization of the Creatophone [...] would be designed from a number of loudspeaker, amplifier, and mixer components especially selected for this use. A set of 16 pairs of loudspeakers and 16 stereo amplifiers would be deployed. In order to keep costs down and make the task of spatialisation manageable in concert, the mixer would be limited to 8 output channels. This means that certain loudspeaker pairs would be grouped, that is, assigned to a single potentiometer. [Note the references to *pairs* of loudspeakers and stereo amplifiers: we will return to this later.][202]

The Creatophone also 'take[s] advantage of the particularities of certain loudspeakers,'[203] and as such is comparable to the Acousmonium (described above) and the Cybernéphone (to be described below).

Roads *et al* also plan to develop a system for the simultaneous diffusion of multiple independent coherent audio source sets, additionally (and

[202] Roads, Kuchera-Morin and Pope (2001). "The Creatophone Sound Spatialisation Project". Website available at: http://www.ccmrc.ucsb.edu/wp/SpatialSnd.2.pdf.
[203] Ibid.

correctly) observing that 'one-fader-to-one-loudspeaker' control would be insufficient for this kind of performance:

> As the project evolves the Creatophone can be directed towards research in control of spatialisation in concert from multiple input sources. This [...] requires computer-controlled digital mixing technology.[204]

It is unclear – owing largely to a lack of detailed published material – exactly how many of the development objectives for this system have actually been realised at this time. When used for the diffusion of electroacoustic works at the Sound in Space Symposium – held in March 2000 at CREATE – the system was described by Harley as follows:

> The Creatophone is a spatial sound projection system of flexible configuration. For this symposium, the various inputs (ADAT, DAT, CD, computer with 8-channel output) were fed into a 16-bus Soundcraft mixer and then out through four Threshold stereo power amplifiers via Horizon, AudioQuest, and Tara interconnects and MIT and AudioQuest speaker cables to eight B&W Matrix 801 loudspeakers. [...] The loudspeakers were placed in a standard octagonal configuration with two in front, two in rear, and two each on either side. [It can be deduced that this array is somewhere in between the two configurations illustrated in Figure 11 and Figure 12 (page 136). Such arrays are – as far as the author understands – fairly standard in the United States.] The best listening was found in the center, as one would expect, but many spatialized effects could be heard quite effectively from other locations as well. The superior clarity and power of the Creatophone sound system certainly enhanced the listening experience, regardless of seating.[205]

Clearly this description falls short of the developers' 'first full realization of the Creatophone' (as described in the quotation given previously) and is presumably not, therefore, representative of the complete system. As little additional information has been forthcoming, evaluation of this system (as far as this is possible) will have to be on the basis of the proposed developments described in the literature.

4.10.2. Evaluation

The Creatophone system as described by Harley is, in fact, no different from a generic sound diffusion system and, according to the quotation given above, would appear to fall into the category of GDS 1 (see

[204] Ibid.
[205] Harley (2000). "Sound in Space 2000 Symposium". *Computer Music Journal*, 24(4)

section 4.6.1), as no splitter box is mentioned and the system comprises only a small number of loudspeakers. Having said that, this information is – at the time of writing – more than five years old, and it can be presumed that subsequent developments may have taken place during this time.

The developers express an interest in 'the peculiarities of certain loudspeakers,' which strongly suggests that loudspeakers that 'colour' the sound are to be utilised. The design objectives also seem to favour stereophonic audio sources, insofar as there is no mention of multichannel sources but much reference to 'loudspeaker pairs,' 'stereo amplifiers,' and so on. Furthermore, simultaneous control of multiple loudspeakers with a single fader has been described as a possibility. This would, of course, be beneficial from the standpoint of interface ergonomics. On this basis, this system appears to be rather similar in its objectives to the Acousmonium, although almost certainly less biased in favour of top-down methods.

Perhaps most interestingly, the CREATE team express an interest in developing tools for the simultaneous independent diffusion of multiple CASSs. This is an area of research that the author believes would be of tremendous benefit within the field of sound diffusion in terms of accommodating works from across the creative and technological spectrum of the electroacoustic idiom, not least because none of the systems reviewed easily facilitate the real-time realisation of such practice. Similar interests are expressed by Harrison and Truax (see section 3.5), although fully developed systems are not yet in evidence.

With these proposed developments borne in mind it would appear that the Creatophone might be of considerable interest and benefit to top-down practitioners, and accordingly perhaps less useful to those with more characteristically bottom-up concerns.

4.11. The Cybernéphone

4.11.1. Description

The Cybernéphone – known up until 1997 as the Gmebaphone – has been under continuous development at the *Institut de Musique Électroacoustique de Bourges* (IMEB) since 1973. The sixth and latest version of the system was first used in performance at the 24th *Concours International de Musique et d'Art Sonore Électroacoustique*, in Bourges, in 1997, and is described at length by Clozier.[206]

Fixed medium audio sources are transferred onto the hard drive of a computer. In performance the computer mixes the audio source signals, played back directly from the hard drive, to fifty loudspeaker output channels, according to input obtained from an intermediate user interface. According to Clozier, there are two primary user interfacing methods: 'manual mode' and 'computer-assisted diffusion mode.'[207] In 'manual mode,' the performer effectively has direct one-to-one control over individual loudspeaker output levels via faders, in much the same way as with BEAST and the 'generic' diffusion systems. 'Computer-assisted diffusion mode' essentially allows predetermined fader movements to be triggered during performance. Such fader movements can be pre-programmed in two ways: either they can be 'recorded' directly from the faders into a sequencer; or parameters can be specified offline in 'configuration tables.' Pre-programmed actions can then be triggered manually during performance, or automatically by a sequencer.

In addition to simple mixing between output channels in real time or semi-automatically, the Cybernéphone also allows the diffuser to perform various treatments on the sonic material: these include 'time,

---

206 Clozier (1998). "Presentation of the Gmebaphone Concept and the Cybernephone Instrument". In Barrière and Bennett [Eds.] *Composition / Diffusion in Electroacoustic Music*: 266-281.
207 Ibid.: 269.

timbre, space, phase and detuning.'[208] Although Clozier does not describe exactly what these processes involve, one might deduce that the system allows the performer to perform time-delay, filtering, amplitude diffusion, phase adjustments and pitch shifting on the audio signals. Again, these actions can be performed in real time or sequenced.

An important aspect of the Cybernéphone that differentiates it from many other diffusion systems is the loudspeaker configuration used. Notably, many of the loudspeakers used have limited frequency response bands. Consisting of a number of such loudspeakers, these 'V-systems,' as Clozier describes them, 'analyze and select the timbres and redistribute them in six sound color registers for each of the left and right channels (2 basses, 2 mediums, 2 trebles) which are sent to special loudspeakers. When the latter are set up, from the very lowest bass to the very highest treble, they effect an acoustic resynthesis of the sounds.'[209] In other words, discrete bands of the audio spectrum are reproduced by loudspeakers with different frequency responses (essentially 'filter characteristics') in different locations, thus distributing the sound in space differently according to its frequency content. The 'bandpass filter' characteristics of each V-system are such that, across the group of loudspeakers, the entire frequency spectrum is represented, but not by any one, single, loudspeaker. In addition to these 'registered' groups of loudspeakers, 'reference' loudspeakers are also provided. These exhibit a linear frequency response across the audio spectrum but can be subject to the user-specified signal processing procedures mentioned in the previous paragraph. Loudspeakers are arranged in pairs, symmetrically about the front-to-back axis of the hall, but are conceptually sub-grouped in various combinations into 'planes' and 'diagonals,' again foreshadowing the concept of the coherent loudspeaker set.

208 Ibid.: 268.
209 Ibid.: 268.

4.11.2. Evaluation

Clozier describes the Cybernéphone as follows:

> The Cybernéphone may be defined as a huge acoustic synthesizer, an interpretation instrument that the composer plays in concert, an instrument that serves to express his composition, to enhance its structure for the benefit of the audience, to bring it to sonic concretization.[210]

Referring back to section 3.12.1 in the previous chapter, it could be deduced that the Cybernéphone – at least, as conceived by Clozier – is very much oriented towards top-down diffusion techniques, in-keeping with the idea that a diffusion system should offer the performer as much interpretative scope as possible, rather than being focused on providing an essentially 'neutral' reproduction of the electroacoustic work. Nonetheless, in terms of the techniques and processes offered to the performer, the Cybernéphone in fact caters for both top-down and bottom-up methodologies extremely well and without obviously exhibiting any bias toward one or the other.

The two parallel modes of operation – 'manual mode' and 'computer assisted diffusion mode' – would seem immediately and obviously conducive to top-down and bottom-up approaches to sound diffusion, respectively. In manual mode, diffusion can be executed in real-time, in a manner that is perceptually effective and improvisatory as far as this is appropriate to the piece and the performance context; this is firmly in-keeping with the top-down aesthetic. In computer assisted diffusion mode, on the other hand, actions that have been pre-programmed offline can be executed automatically. This is more in tune with the bottom-up aesthetic because it allows the parameters of diffusion actions to be precisely specified outside of the performance context (this can, accordingly, be done with respect to any kind of deterministic scheme as required, as opposed to being 'whimsically improvised' at the moment of performance) and triggered at an objectively defined and mathematically accurate time.

[210] Ibid.: 268.

Another advantage is that there would seem to be significant facility for cross-over between these two fairly 'polar' modes of operation. Predetermined diffusion actions, for example, can either be specified offline as described above, or else recorded in real-time via the control interface. Although the offline option would seem more immediately affiliated with the bottom-up approach, it could certainly be useful to the top-down performer as a means to facilitating actions that would be difficult or impossible to attain in real time, for example. Equally, a slightly more 'liberal' bottom-up diffuser may wish to define diffusion actions in real time, but have those real time actions consistently and accurately reproduced throughout the performance (perhaps also wishing to achieve a degree of uniformity across several different performances of the same work). In addition, these predetermined actions – howsoever obtained – can either be triggered manually in real-time, or automatically at a specified time, facilitating a further dimension of 'crosstalk' between top-down and bottom-up techniques.

In comparison with those systems described earlier, it can also be observed that the Cybernéphone has rather more scope in terms of the *range* of diffusion actions available. By and large, the systems accounted previously only allow the performer to manipulate the relative amplitudes of various input-to-output signal routings,[211] whereas the Cybernéphone – in addition to the usual amplitude diffusion – offers time-delay, filtering, phase adjustment and pitch shifting capabilities. Needless to say, each of these could be utilised in either an essentially top-down ('creative,' in the sense of continuing the compositional process) or bottom-up (perhaps more 'corrective,' directed toward achieving a more transparent presentation of the finished work) manner and their inclusion is therefore extremely beneficial from both perspectives. The benefits are augmented in that

[211] Although certain systems – particularly those based around conventional studio mixing desks – may offer certain techniques including EQ and phase inversion, for example, in most cases these cannot be ergonomically controlled to the same extent possible with the Cybernéphone. Of course, such things could be implemented in experimental approaches such as those described in section 4.9.

these processes can be recorded and automated in each of the ways described previously.

In terms of signal routing and interface ergonomics, it is implied that loudspeakers can be grouped together to a certain extent and controlled with a single 'fader.' This suggests a higher degree of fulfilment of criteria 4.3.3 and 4.3.4 than that attained by those systems operating on a one-fader-to-one-loudspeaker basis, although the literature does not go into much detail in these respects. It is not particularly clear, for example, exactly *how* flexible the signal routing and mixing architecture is, nor to what extent the grouping of loudspeakers is fixed or dynamic. Certainly, the presence of both 'registered' (coloured) and 'reference' (transparent) loudspeakers – particularly in conjunction with the ability to dynamically apply equalisation – would be beneficial to both top-down and bottom-up practitioners, depending of course on the extent to which the diffuser has control over which loudspeakers are used at any given time. Again, this is not particularly clear in the sources consulted.

Like all of the systems described so far, the Cybernéphone is based around a control interface that is essentially comparable to a conventional studio mixing desk and – like BEAST's experimental system – this is abstracted from the mix engine via the MIDI protocol, adding a further dimension of flexibility. Personal communications with the author suggest that this interface is accessible as per criterion 4.3.5 (according to my sources, it does not 'feel' enormously different from a standard mixing desk) although the touch-sensitive strips are reported to respond less well with sweaty hands, which can be problematic in the stressful (and often rather warm) performance context.[212]

Moore *et al* observe that the Cybernéphone 'is only suited to the diffusion of stereo works from CD or DAT and is very difficult to adapt

[212] Adrian Moore (2002) and Nikos Stavropoulos (2002). Personal communications with the author.

to other formats.'[213] This seriously restricts the fulfilment of criterion 4.3.1, as works with different technological demands cannot be accommodated. However, unlike any of the other systems to be described, the Cybernéphone addresses the issue of diffusion 'practice' outside the performance context:

> The entire program [software] is made available on CD-ROM, allowing the composer to prepare his diffusion and record it in sequence. The diffusion can then be played as is, adapted to the conditions in the concert hall and to the audience, or it may be viewed as a temporary configuration, a guide sequence that can be interpreted in terms of dynamic timbres and spaces, all in real-time, by a musician at the console. Second, the composer can send his diffusion program on CD-ROM to be used in concert in Bourges or elsewhere, or he can carry out the diffusion himself via Internet.[214]

It is obvious that both top-down and bottom-up methodologies have been considered, and that in allowing performances to be prepared elsewhere the important issue of rehearsal time and familiarity with the instrument has been addressed.

In summary, the Cybernéphone caters very well for both top-down and bottom-up approaches to sound diffusion by offering an extended range of diffusion actions (in addition to the standard amplitude diffusion) that can have useful applications from both perspectives. It also benefits from a capability that has not been observed thus far: automation. Although this facility would, at first, appear more beneficial from the bottom-up perspective, it would appear that this particular implementation also has demonstrable benefits for the top-down diffuser – as described previously – although the literature available does not confirm or negate this deduction as assuredly as would be ideal.

---

[213] Moore, Moore and Mooney (2004). "M2 Diffusion - The Live Diffusion of Sound in Space". *Proceedings of the International Computer Music Conference (ICMC) 2004*
[214] Clozier (1998). "Presentation of the Gmebaphone Concept and the Cybernephone Instrument". In Barrière and Bennett [Eds.] *Composition / Diffusion in Electroacoustic Music*: 267.

4.12. Sound Cupolas and The Kinephone

4.12.1. Description

It is interesting to compare the stipulation, as proposed above by Clozier and by many others, that an electroacoustic work is 'concretised' at the performance stage, via a process that essentially continues the top-down method of working, with the following, highly contrasting, assertion:

> The spatial projection of a composition is realized first and foremost in the studio, where the composition progressively takes shape according to a precise working method... If the serious listening public is placed in a manner unfavourable to the perception of the intended spatial composition, then the spatial effect will be lost, and the composer, who has spent many weeks at the difficult task of creating clear and precise spatial forms will have wasted his time and energy... The studio in which composition takes place, and of necessity the concert hall, must be as anechoic as possible if there is not to be a serious risk of confusion of parameters.[215]

It is fundamental musical beliefs such as these that have informed the development of 'sound cupolas' – purpose-built standardised listening spaces – from 1977 onwards, and sound diffusion systems (in some respects) radically different from those described so far.

A sound cupola essentially consists of a hemispherical or hemi-ovoid frame – erected within a performance venue – to which a large number of loudspeakers are attached. All of the loudspeakers are angled towards the central, focal, point of the symmetrical frame 'and then very exactly adjusted so as to obtain equal intensity, phase and spectrum.'[216] The natural acoustic characteristics of the venue are typically suppressed by covering the top of the frame structure with cloth, or some other acoustically absorbent material:

> The structures and bridges of the scaffolding conceal all the cables, and the cloth covering insulates the hall and renders it nearly anechoic.[217]

[215] Küpper (1998). "Analysis of the Spatial Parameter: Psychoacoustic Measurements in Sound Cupolas". In Barrière and Bennett [Eds.] *Composition / Diffusion in Electroacoustic Music*: 289.
[216] Ibid.: 296.
[217] Ibid.: 296

The number and positioning of loudspeakers in sound cupolas is determined by research into the acoustics and psychoacoustics of directional sound localisation. Küpper describes a series of experiments undertaken in order to determine 'the spatial interval,' that is, a notionally 'universal' subdivision of three-dimensional auditory space into discrete intervals measured in degrees. Comparison with the processes of 'compartmentalisation' described in section 2.2.2 should identify such practice as a characteristically bottom-up approach. Further to this research, Küpper proposes that the number of spatially discriminable points on a notional 'auditory sphere' surrounding the listener, is somewhere in the region of six-thousand; that is, the supposed 'resolution' of the human hearing system, in terms of spatial localisation, allows us to perceptually differentiate between approximately six-thousand discrete sound source locations. In between these points, differentiation is less reliable or impossible. Küpper also notes that, 'only 88 out of a maximum of 1400 determinable pitches were chosen for the Western tempered scale (12 pitches per octave) for a ratio of 1 out of 15.9.' Using the same ratio, Küpper eventually states that the spatial equivalent of a semitone:

> ...varies between 13.3 degrees and 15.8 degrees. We therefore propose 15 degrees as the tempered spatial interval, since 360 degrees may be divided easily by this figure.[218]

Accordingly, loudspeakers in sound cupolas are, as far as possible, positioned a 'spatial semitone' (15 degrees) apart, ultimately resulting in around 102 equally spaced loudspeakers for a hemispherical or similarly shaped array. A sound cupola installed at Linz in 1984 consisted of 104 loudspeakers attached to a hemi-ovoid frame. This configuration was used for the performance of fixed medium electroacoustic works, as well as pieces featuring live sound encoded via microphones.

[218] Ibid.: 294.

In line with the bottom-up ethos, Küpper implies a strong distaste for overly improvised, whimsical, or 'intuitive' diffusion, preferring more 'objective' schemes, or even automated computer realisations:

> Four interwoven spirals [...] enable the performers (one, two or four in number) to set up spatial play more easily. Based on these four very symmetrical spirals, the performers play tape compositions as well as live concerts. In the light of previous experience, these four spirals are a safety factor in helping to avoid clumsy, improvised or chance spatial movements. Spatial performers do not yet exist, and in the future it would be good if composers themselves were able to construct the spatial structures and movements on computer.[219]

When Küpper states that 'spatial performers do not yet exist,' what he means is that there is no universally accepted objective 'scale' for the spatialisation of sound within which a performer can exhibit virtuosity. In this context, any diffusion that is based 'merely' on subjective judgements – without the need for, say, a 'spatial scale,' or a 'score' – is completely discredited on the grounds that there is no objective milestone against which to judge the quality or 'correctness' of the performance, an attitude to which many top-down diffusers would surely (and quite justifiably) object. It is also interesting to note that Küpper describes the four predetermined spatial trajectories as 'a safety factor in helping to avoid clumsy, improvised, or chance spatial movements.' This implies that the spiral trajectories could – in Küpper's opinion – be meaningfully applied to any musical material, regardless of the perceived nature of that material, an assertion that is indicative of an absolutely characteristic bottom-up attitude. Firstly, it suggests that deterministic musical structures can be (and, indeed, *should* be) imposed upon 'fundamentally inert' musical material (recalling Harrison's quotation, cited in section 3.13): it therefore 'does not matter' what the perceived relationship between the abstract structure and the musical material is; this is simply not a consideration. Secondly, in accordance with this, the proposition is that the spatialisation (via carefully planned diffusion) adds a previously non-existent dimension of interest and structural rigour to the musical work, 'where once there was nothing' (recalling observations made in section 2.5.1, further to

[219] Ibid.: 296.

Savouret's writings). Again, this position simply does not recognise the belief that musical materials might have spatial characteristics perceptually 'embedded.' Top-down practitioners would tend to disagree strongly with this position, as anecdotally accounted by Harrison:

> I went to a conference in the States, and the guy on the desk 'diffused' [...] all the pieces. And what he actually did was, he had choreographed, on the mixer, certain fader movements, and he did them irrespective of what material was coming off the medium. So he just went through the same routine: he just went from these faders, to these faders, from these faders to these faders, to the next, the next, and then he started again. And he did it at the same tempo, he did the same pattern, the same sequence. He was just arbitrarily moving the sound, no matter how fast it was coming, no matter whether it was lots of movement implied or no movement implied, he was moving it round the room arbitrarily, in the same sequence, piece after piece after piece. That is not diffusion: that's ridiculous; that's just nonsense.[220]

Whether or not this is in fact 'nonsense' ultimately depends on one's musical beliefs. In the context of top-down ideas the situation does, indeed, appear to be completely nonsensical. To a bottom-up practitioner, however, the same approach may seem perfectly legitimate, particularly in comparison with a (perhaps badly) improvised performance. In fairness, the example scenario is almost certainly not representative of a high quality performance (neither top-down nor bottom-up), and is probably fully deserving of criticism, but it has nonetheless served to articulate an important point.

4.12.1.1. The Kinephone

Perhaps developed in response to the belief that 'spatial performers do not yet exist,' the Kinephone is a performance interface, based on the piano keyboard, that allows a diffuser to control sound diffusion via a large array of loudspeakers. The specific instrument described by Küpper has fifty keys, each one of which, when pressed, opens a particular loudspeaker channel. It was used in conjunction with a cupola installed at the *Palazzo*

[220] Harrison (2004). "An Interview with Professor Jonty Harrison". Personal interview with the author, 8th September 2004. See Appendix 1 for full transcription.

Sagredo, Venice, in 1986.[221] Such an interface would seem particularly appropriate for a diffusion system whose loudspeakers are positioned according to the theory of the 'spatial semitone.' As Küpper also points out, it allows diffusions to be scored using traditional western notation, thus ensuring that the performance is not improvised. (Recall that Piché, cited in section 3.12.2, also called for diffusion practice to be 'codified').

4.12.2. Evaluation

It should already be evident from the previous section that sound cupolas and the Kinephone have been devised according to a fairly puritanical bottom-up agenda. It is therefore reasonable to suppose that the system will be intrinsically biased towards bottom-up diffusion methods. Of course most systems can be appropriated in a broadly top-down or bottom-up manner depending on the specific approach of the diffuser in question, but it is nonetheless clear that this particular system has been designed almost exclusively with respect to bottom-up criteria.

One fairly unusual aspect is that this system offers three-dimensional sound source diffusion, with height. Although loudspeakers positioned above (and perhaps even below) the audience are not entirely uncommon,[222] it is generally true to say that the most frequently employed loudspeaker arrays exhibit an enormous bias in favour of the horizontal plane. As such, diffusion in sound cupolas engenders certain creative possibilities that are rarely available elsewhere.

Use of the Kinephone to control signal routings to large numbers of loudspeakers raises the very important issue of interface ergonomics. In the earlier sound cupolas, control of signal levels sent to loudspeakers

[221] Küpper (1998). "Analysis of the Spatial Parameter: Psychoacoustic Measurements in Sound Cupolas". In Barrière and Bennett [Eds.] *Composition / Diffusion in Electroacoustic Music*: 296-297.
[222] The system installed at the Sonic Arts Research Centre (SARC) in Belfast, for instance, has loudspeakers positioned above and below the audience. This system will very briefly be described in section 4.15.

was – as in several of the systems described previously – via a fader-based mixing desk control interface, with levels being controlled on a one-fader-to-one-loudspeaker basis. With up to 104 loudspeakers, four diffusers were required. The Kinephone undoubtedly represents a more ergonomic design for this particular purpose, but only if the performer is a reasonably accomplished pianist! In this respect, the Kinephone – as a control interface – is not particularly accessible as per criterion 4.3.5 (particularly if an accurate 'bottom-up-style' diffusion is sought; this is, after all, the *forte* of this particular system) and it could be argued that much practice would be required to master it. Nonetheless it could also be argued that the piano keyboard is, at least, a *comparatively* familiar interfacing paradigm, although its suitability for the purpose of sound diffusion would remain open to debate.

The evidently non-portable nature (criterion 4.3.7) of sound cupolas also raises issues, particularly in conjunction with the unusual nature of the Kinephone interface. With more than one hundred loudspeakers it seems extremely unlikely that a cupola could be installed in anything other than a sizeable venue. In addition to this, there are quite considerable financial implications and practical constraints regarding the construction of such systems. These observations indicate that the present diffusion system is also highly inaccessible as per criterion 4.3.6, and therefore only likely to be of any real use to those fortunate enough to have regular and easy access to it.

Overall, sound cupolas and the Kinephone are effectively completely biased towards bottom-up diffusion methodologies, and are – realistically – highly impractical and inaccessible. However, such systems do give an excellent indication of kinds of requirements demanded by classically bottom-up practitioners, and this will surely aid the design of future sound diffusion systems.

4.13. The ACAT Dome

4.13.1. Description

Comparable in some ways to Küpper's sound cupolas, the ACAT Dome is a scaffold geodesic dome measuring fourteen metres in diameter. Sixteen loudspeakers are symmetrically arranged on the frame, as well as five cinema screens for visual image projection. Sound is distributed among the sixteen loudspeakers using the Ambisonic system via a software application running on a Macintosh computer.

Some highly relevant advantages of the Ambisonic system will be cited later and on this basis an explanation is necessary. Ambisonics allows any location in three-dimensional auditory space to be represented via what is effectively an abstraction of a coherent audio source set, comprising four discrete encoded audio streams.[223] These signals – collectively known as Ambisonic B-Format and individually referred to as *W*, *X*, *Y*, and *Z* – are rather like the axes on a three-dimensional Cartesian graph, to use an intuitive visual analogy. The amplitude of signal *X* in relation to the others, at any given time, denotes the position of the sound source on a 'left-to-right' axis, with the relative amplitudes of *Y* and *Z* respectively denoting the position on 'front-to-back' and 'up-to-down' axes. The amplitude of *W* does not change, and this is basically to improve the quality of the perceived effect upon reproduction. Monophonic signals can be positioned within an Ambisonic 'sound field' by substituting polar co-ordinates (azimuth, elevation) relating to the desired three-dimensional location into a set of four mathematical equations (which, when resolved, result in each of the four constituent B-Format channels), or else recorded with a Soundfield microphone comprising four capsules, and an encoding process that derives the B-Format signals from the microphone capsule signals. Very importantly, it should be noted that unlike a coherent

[223] First-order Ambisonics comprises four signals. Higher-order equivalents involve more channels but the principles remain the same, and in any case, it is not necessary to discuss these in this context.

audio source set, Ambisonic B-Format is not composed of encoded audio streams that can be fed directly to loudspeakers. The B-Format signals must first be decoded into loudspeaker feed signals – again, mathematically – in order for the desired effect to be achieved. Because B-Format is an abstract representation of a three-dimensional sound field, so it can be decoded for a variety of different loudspeaker arrays (using formulae incorporating the relative positions of the loudspeakers) ostensibly with the same spatio-auditory results. Obviously, the quality of the final results depends on the number of loudspeakers, and it is generally accepted that three-dimensional auditory results are not convincingly attainable with less than eight loudspeakers arranged in a cubic formation.

Returning to the diffusion system in question, mouse movements are used to control real-time sound source positioning on the horizontal plane, while a variable height foot pedal provides the vertical positional information. Data is communicated to the software via the MIDI protocol. The system, developed at the Australian Centre for the Arts and Technology (ACAT), is described more fully by Vennonen.[224]

4.13.2. Evaluation

A fairly unique characteristic of this system rests in the fact that it uses the Ambisonic system, which has an altogether different set of operational affordances and applications in comparison with those systems described previously. Abstracted control is a particularly useful characteristic; this will be discussed more fully in Chapter 6. A detailed discussion of Ambisonics is beyond the scope of the present thesis, but the salient aspects of the system are summarised by its inventor, Michael Gerzon,[225] and also by the author in a previous thesis.[226]

[224] Vennonen (1994). "A Practical System for Three-Dimensional Sound Projection". Website available at: http://online.anu.edu.au/ITA/ACAT/Ambisonic/3S.94.abstract.html.
[225] The standard primary source for Ambisonics is:

Gerzon (1974). "Surround-Sound Psychoacoustics". *Wireless World*, **80**(12).

A useful bibliography of other relevant sources, compiled by Gerzon, is:

Gerzon (1995). "Papers on Ambisonics and Related Topics". Website available at: http://members.tripod.com/martin_leese/Ambisonic/biblio.txt.

The computer mouse is a control paradigm that is familiar to most people nowadays, and in this case it is used in relation to spatial positioning in a two-dimensional plane, which is not too far removed from its normal application. It seems likely, therefore, that this operational paradigm could be assimilated by performers with relative ease (in comparison with, say, the piano-style interface of the Kinephone, which is potentially far less approachable). The efficacy of the foot-pedal is slightly more difficult to evaluate as this control paradigm – although not by any means uncommon – has a wider variety of potential applications.[227] Certainly, the simultaneous operation of a mouse and foot-pedal presents certain challenges with respect to motor coordination, and on this basis it seems reasonable to suppose that the control interface might take some time to get to grips with. Therefore, the overall efficacy of the system will be determined to a certain extent by its widespread availability for practice purposes (relating to criterion 4.3.6), and in this respect, the ACAT Dome is open to some of the same criticisms directed towards Küpper's system, described in the previous section.

Additionally, it is not immediately apparent (in terms of interface ergonomics) how a single mouse and foot-pedal could be used to simultaneously diffuse multiple sources, nor as a means to realising pluriphonic diffusion. This system effectively allows one CASS at a time to be positioned as a 'point source' and it can therefore be argued that the system is perhaps more suited to 'panning' than sound diffusion (see section 3.10). It might therefore be suggested that this system is biased in favour of the bottom-up ethos whereby 'space' is performed live as an additional and autonomous dimension to the electroacoustic work (in this context the distinction between 'panning' and 'diffusion'

---

[226] Mooney (2001). "Ambipan and Ambidec: Towards a Suite of VST Plugins with Graphical User Interface for Positioning Sound Sources within an Ambisonic Sound-Field in Real Time". M.Sc. thesis (University of York).

[227] Nonetheless, in abstract terms the foot-pedal is – in the author's opinion – a useful addition to the diffusion repertoire because it adds a further dimension of interactive control. It can therefore be argued that time spent assimilating the necessary skills is, at least potentially, worthwhile.

is less obvious, in any case). In this respect, the ACAT Dome could be considered similar – in terms of its ethos – to Küpper's system, described in the previous section. Indeed, the Ambisonic system – with its emphasis on the accurate and realistic reproduction of three-dimensional auditory phenomena – could itself be considered biased in favour of bottom-up ideologies.

4.14. The DM-8 and the Richmond AudioBox

4.14.1. Description

The DM-8 and its various subsequent developments represent diffusion systems built upon the concept of a *mix matrix*, which is described by Rolfe as follows:

> A matrix mixer differs from a conventional mixer in that the usual on/off buttons in the assign-to-buss section are replaced by level controls. An 8 input, 8 output mixer has, for example, 64 (8 times 8) independently adjustable levels commonly called *crosspoints* and usually conceived of as a square, or *matrix*.[228]

The 8x8 matrix to which Rolfe refers is illustrated in Figure 29, below.

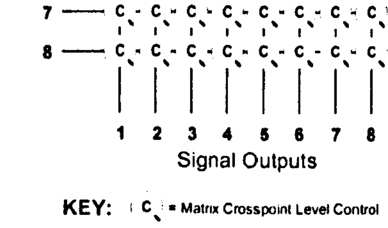

Figure 29. Illustration of an 8x8 signal routing matrix with a level control at each crosspoint.

228 Rolfe (1999). "A Practical Guide to Diffusion". *EContact!* 2(4). Electronic journal available at: http://cec.concordia.ca/econtact/Diffusion/pracdiff.htm.

Crosspoints can be referred to like grid references on a map: a horizontal reference followed by a vertical reference. For example, if we turn up the level on crosspoint 1-1, we are mixing some of the audio at signal input 1 to signal output 1. If the controls at crosspoints 1-1, 1-2, 1-3 and 1-4 are all turned up, then the signals present at inputs 1, 2, 3 and 4 will be mixed together at signal output 1, with the balance of the mix proportionate to the relative positions of each of the crosspoint controls in question.

It should be immediately apparent that this model is potentially far more flexible than, say, the routing matrix implemented in BEAST's DACS 3D mixing console, which only allows routings to be switched on or off. Further, unlike most switch-based audio mixing hardware, the mix matrix paradigm is not inherently biased towards stereophonic coherent audio source sets insofar as it allows inputs to be mixed to outputs arbitrarily. However, performing live sound diffusion using a bank of sixty-four rotary potentiometers as illustrated in Figure 29 is, realistically, impossible. This brings us back to the issue of interface ergonomics. Accordingly – like the Cybernéphone – the DM-8 and those systems developed from it represent diffusion systems based on both software and hardware components. In other words, control of the matrix crosspoints is not on a physical, one-by-one basis, but achieved via an intermediate software interface that allows crosspoint level controls to be indirectly manipulated in a way that the limits of human dexterity would otherwise restrict.

In the DM-8 system, the mix matrix itself is realised in hardware, and is controlled via a software interface written in Max/MSP. The interface is described in some detail by Rolfe and is 'speaker-centric,'[229] that is, predetermined diffusions are specified in terms of which loudspeakers they utilise at which times. This means that the loudspeaker array used for the performance of the predetermined diffusion must be identical to the array used to create the diffusion. This problem was addressed in a

[229] Ibid.

subsequent development of the system comprising a 16x16 hardware matrix mixer (namely the Richmond AudioBox, marketed by Harmonic Functions[230]) and a Max/MSP interface, written by Rolfe and others, named *ABControl*.[231] With this system – again, described in more detail by Rolfe – crosspoint levels themselves are determined 'on the fly' via an intermediate vector-based user interface:

> A rotation [for example] mapped onto vector angles can then be played back on a quad, 8-channel, or other configuration.[232]

Nonetheless, one assumes that the loudspeakers would still have to be arranged in a circular formation around the audience for the desired effect to be achieved.

4.14.2. Evaluation

This system benefits enormously by offering performers the ability to express diffusion actions abstractly. For example, a rotation around the loudspeaker array can be expressed in terms of the parameters of the rotation itself (rate of rotation, direction of rotation, *et cetera*), as opposed to the sequence of mix matrix parameter adjustments required to achieve the same effect. Of the systems described thus far, only the ACAT Dome offers this kind of functionality, the other systems all effectively operating directly on loudspeaker amplitude settings. This is because the control interface (in this case realised entirely in software) and mix engine (in this case the hardware AudioBox matrix mixer) are not directly linked, but rather separated by an intermediate abstraction layer in which control data can be algorithmically transformed into the data required by the mix engine.[233]

[230] Harmonic Functions (2001). "The AudioBox Disk Playback Matrix Mixer". Website available at: http://www.hfi.com/dm16.htm.
[231] Rolfe (1999). "A Practical Guide to Diffusion". *EContact!* 2(4). Electronic journal available at: http://cec.concordia.ca/econtact/Diffusion/pracdiff.htm.
[232] Ibid.
[233] Although several of the systems described previously (namely, the Acousmonium, BEAST's experimental system, Cybernéphone) do, in fact, operate on this basis, these systems opt for control paradigms that closely emulate direct parameter control.

A fundamental characteristic the DM-8, and of those systems developed from it, is that these are completely offline, automated, systems. In other words, the diffusion must be planned and programmed in advance for strictly automatic playback during performance. In terms of their ethos, this essentially limits the application of these systems to use by bottom-up practitioners. Indeed, the DM-8 might not even be regarded as diffusion system at all by top-down practitioners, as it does not allow for the real-time shaping of musical material in a way that is perceptually effective for the audience – a criterion upon which sound diffusion, from a top-down perspective, is fundamentally based. In order to be appropriate for top-down use, the diffusion would have to be realised in the performance venue itself, and this will not normally be possible due to insufficient time being available. In any case, the acoustic characteristics of the venue are likely to be considerably different when the audience is present, and this can – of course – completely change the diffusion requirements, effectively rendering a pre-planned performance useless. This system is therefore not particularly useful for the performance of top-down works, and consequently falls short of fully satisfying criterion 4.3.2.

Because the DM-8 is an offline system, so considerations with respect to its interface ergonomics will have a rather different emphasis. In this case the interface is realised entirely in software and based around standard keyboard and mouse input devices. As observed in section 4.13.2, these are not tools best suited to real-time sound diffusion but are, however, extremely well supported in the context of offline human-computer interaction. Accordingly, systems based around the AudioBox have been reasonably widely employed as *compositional* tools, the offline editing paradigm being particularly well-suited for multichannel panning, which can be difficult using more mainstream audio sequencers. Such a system is currently in use, for example, in the electroacoustic music studios at Leeds University.[234] The degree of cross-over between 'compositional' tools and 'performance' tools

[234] Ewan Stefani, University of Leeds (2004). Personal communication with the author.

apparent in electroacoustic music has already been noted and is something that should be considered in the design of new sound diffusion systems.

4.15. Other Systems

The Sonic Arts Research Centre (SARC) at Queen's University, Belfast, has a forty-eight channel sound diffusion system permanently installed in its sonic laboratory.[235] Loudspeakers are deployed at various heights in four 'tiers,' one of which is underneath the acoustically transparent grid floor of the listening auditorium. Diffusion is controlled with a Digidesign Control-24 mixing surface, meaning that the SARC system effectively operates in the same way as a generic diffusion system.

The elBo, devised by Colby Lieder at Princeton University, is described is follows:

> The elBo is an instrument for live spatialization and diffusion of sound in space. The instrument translates gestural actions into spatial location, point source velocity, pan width (spread), effects parameters, and other diffusion data. The guts fit in a single rack unit chassis.[236]

What Lieder describes is not a full 'diffusion system' so much as a control interface that could conceivably be incorporated into one. Very little other published information is available.

Along the same lines is a device known as the aXi0, which was 'designed and built by Brad Cariou at the University of Calgary and conceived to provide digital artists with the intimate control and flexibility needed to express themselves clearly in various new media.'[237] It consists of a velocity-sensitive MIDI keyboard, a joystick, and an array of buttons; collectively, these facilitate multidimensional control over parameters. The precise nature of the control is determined by the user in a Max/MSP interface, so the device is not limited solely to sound diffusion applications.

---

235 SARC (2005). "Sonic Lab". Website available at: http://www.sarc.qub.ac.uk/main.php?page=building&bID=1.
236 Lieder (2001). "Colby Lieder - Instruments". Website available at: http://silvertone.princeton.edu/~colby/Instruments.html.
237 Eagle (2001). "The aXi0". Website available at: http://www.ucalgary.ca/~eagle/.

Again, this is a control interface as opposed to a complete sound diffusion system, but represents interesting research into control paradigms that surpass the simple mixing desk fader. However, 'novel' devices such as the Elbo and aXi0 – with particular reference to criterion 4.3.5 – present practitioners with the difficulty of having to appropriate new performance skills. As such, these are more representative of experimental approaches.

There are, of course, many other systems currently in development or in use for the live performance of electroacoustic music. A detailed account of all of them is far beyond the scope of the present thesis, and in any case, there are surely many systems that remain either under-documented or completely undocumented, further complicating the task of achieving anything even close to a comprehensive review.

4.16. Summary

Observations specific to each of the systems evaluated have been given previously. Clearly, the sheer number of different approaches to the task of sound diffusion – as evidenced in the systems described – makes their evaluation on a general level somewhat problematic. The broadest conclusion that can be made is that, at present, no single sound diffusion system that is in regular use for the performance of electroacoustic works is able to adequately satisfy all of the given criteria, and this circumstance can be attributed to a great many factors. The relative successes and failings of each system are variable, but nonetheless, certain general observations can be made.

In technical terms, problems arising from the inflexibilities inherent in mixing hardware – particularly that which has not been designed specifically for the purpose of sound diffusion – are strikingly apparent. If the mix architecture itself limits the ways in which audio input streams can be diffused to audio outputs, this is likely to have a negative impact on both the technological range of works that can be accommodated, and on the range of ways in which diffusion can be realised. This is particularly true if the control interface and mixing architecture are intrinsically linked, a circumstance that effectively compounds the already considerable challenges presented in each of these areas. Such limitations are particularly in evidence in the generic setups, and indeed in all of the systems employing hardware mixing desks (BEAST, Creatophone, Acousmonium). For largely the same reasons, bias in favour of stereo, to the effective exclusion of larger CASSs, is also a problem.

Many of the technical problems can potentially be overcome by the use of software mix engines, as implemented in some of the systems described (BEAST experimental, DM-8, ACAT Dome, Cybernéphone). However, the Cybernéphone and BEAST's experimental system adhere fairly closely to traditional hardware paradigms and, while this is advantageous from the perspective of familiarity, it does not exploit the benefits of software abstraction to their full potential. The DM-8 and ACAT Dome, on the other hand, deviate from familiar paradigms to such an extent that they seriously limit the scope of their applicability. This is particularly unfortunate in the case of the DM-8, because the mix-matrix architecture it employs potentially represents an effective solution to many of the problems discussed. This demonstrates the importance of careful user interface design: a flexible internal architecture is effectively useless without an appropriate control interface. It would be useful to find a 'middle ground' between the relatively conservative emulation of familiar hardware architectures and those paradigms that deviate radically from these.

The subject of interface ergonomics presents interesting challenges, particularly if interfaces that are 'easy to learn' are to be sought. Most of the systems reviewed implement faders or fader-like controls as their primary control interface. In brief, these are familiar (this is an advantage), but present fairly serious problems in terms of ergonomics. Those systems that deviate from this paradigm (Kinephone, ACAT Dome, DM-8), again, limit their applicability fairly seriously. Furthermore, any interface that strays too far from the familiar runs the risk of being inaccessible in the limited amount of rehearsal time available. Again, an approach that mediates between these poles (familiar versus unfamiliar) would be beneficial.

Clearly, none of the systems are as portable as might be desirable. Several of the systems are, realistically, permanent installations, and all have considerable transportation and setup overheads. Again, software systems offer a potentially elegant solution to this problem, but problems arise where systems are operationally dependent on large loudspeaker arrays. A system that would allow similar results to be obtained on a wide variety of arrays for practice purposes would be beneficial from this perspective. A system combining this feature with the availability of diffusion software for independent practice and/or preparation of performances – as demonstrated by the Cybernéphone – would be ideal.

In terms of technological prerequisites, the most readily available systems are the generic setups, and BEAST (whose system has, in line with this observation, been adopted elsewhere). The DM-8 is also fairly easily available, and has been adopted elsewhere. Other systems (the Cybernéphone, for example) rely heavily on custom and/or bespoke components or present considerable logistical implications (ACAT Dome, Sound Cupolas). Unfortunately, the systems that are the most accessible in terms of the technologies required to build them, are not the most flexible.

On a grand scale, it can be observed that, as creative frameworks, the systems described each tend to exhibit a reasonably strong bias in terms of their aesthetic directionality. In other words, they are clearly and easily divisible into those that are bottom-up (Cupolas/Kinephone, ACAT Dome, DM-8) and those that are top-down (BEAST, Acousmonium, Creatophone). This is often a function of the nature of the control interface in particular: interfaces that are amenable to objectively precise control tend to be suited to bottom-up applications, for example. The systems in most common use – the generic setups – are slightly biased in favour of top-down methods (they are best suited to pluriphonic stereo performances and are not particularly conducive to 'accurate' control of the diffusion in objective terms) but, in actual fact, are not particularly flexible from either perspective. The only system that makes clear and obvious efforts to cater for both top-down and bottom-up approaches to diffusion is the Cybernéphone, and much can be learned from this system in this respect.

MISSING PAGE

5. The M2 Sound Diffusion System

5.1. Introduction

The M2 Sound Diffusion System was conceived, designed, developed, and built jointly by Dr. David Moore and the author at the University of Sheffield Sound Studios between 2002 and 2004, hosting its first public performance of electroacoustic works as part of the Electroacoustic Wales concert series at the University of Wales, Bangor, in spring 2004.

While it would be slightly misleading to suggest that the M2 was conceived strictly 'in response' to the research documented in the present thesis (for one thing this is not exactly the case from a purely chronological perspective), it is certainly true to say that the development of the system has played an extremely important role in the research process itself. The prototype design presents many significant improvements in comparison with other diffusion systems and also raises further important issues to be considered in the design of new systems; both of these are highly positive outcomes.

In building a new system one must obviously pose questions relating to the proposed capabilities of that system, and give careful thought to the considerations to be made in its design. It was through posing such questions in the early stages of designing the M2 that the criteria for the evaluation of sound diffusion systems (presented in the previous chapter) began to take shape. The design of a system that would satisfy such criteria to a greater extent than existing diffusion technology was, therefore, an important objective in the development of the M2.

Accordingly, there are several areas in which the M2 clearly exceeds the capabilities of many of the systems reviewed previously. The compact and portable nature of the system, for instance, and the relative speed and ease with which it can be set up, cannot realistically be rivalled by any currently available sound diffusion technology. Careful consideration has also been given to the demands of live sound diffusion from the perspective of signal routing and mixing, an area identified as a deficiency in many other systems. The M2 allows the routing and mixing architecture to be determined by the diffuser, thus offering a degree flexibility unavailable elsewhere. Additionally, the specific actions made available by the custom-designed control interface (or, in very simplistic terms, 'what each fader does') can also be user-defined, a further improvement on the (for the most part) fixed functionalities offered by other systems. Conceived from the outset as a *modular* system, the M2 is also extremely flexible in terms of the technologies required to build it and (with the exception of the software components, to be described presently) does not depend upon specific technologies. This has a number of advantages that will be described in due course.

Such improvements have, of course, been the result of careful planning throughout the design of the system. The development of the M2 has also, however, raised further issues that could not have been predicted by any preliminary research, issues that will clearly be of significance in the design of future systems. In this respect, the relative limitations of the M2, as identified since the completion of the prototype in early 2004, are just as important as its strengths in formulating a useful way forward for the design of new sound diffusion systems. The purpose of the present chapter is, therefore, to present a description and evaluation of the M2 Sound Diffusion System, giving an objective overview of its various strengths and weaknesses. Evaluation will be – as with all of the systems described in the previous chapter – according to the set of criteria outlined in section 4.3.

5.2. System Description

The M2 comprises software and hardware components, the former taking the shape of Moore's *SuperDiffuse* applications.[238] The control interface was designed and assembled by the author, with certain bespoke components being manufactured by the Department of Electrical and

[238] Moore (2004). "Real-Time Sound Spatialization: Software Design and Implementation". Ph.D. thesis (University of Sheffield).

Electronic Engineering at the University of Sheffield to the author's specification.

The control interface consists of a bank of thirty-two high quality ALPS 100-millimetre throw faders arranged in two horizontally-oriented rows of sixteen and housed in a custom designed aluminium box. This is shown in Figure 30. In terms of its functionality, the control interface is completely abstracted from the mix engine, which is realised in software, and communication with the latter is via MIDI controller messages. As such, the control interface does not directly process any audio. Conversion from simple control voltages into MIDI controller messages is achieved via an Ircam AtoMIC Pro interface that has been pre-programmed appropriately.

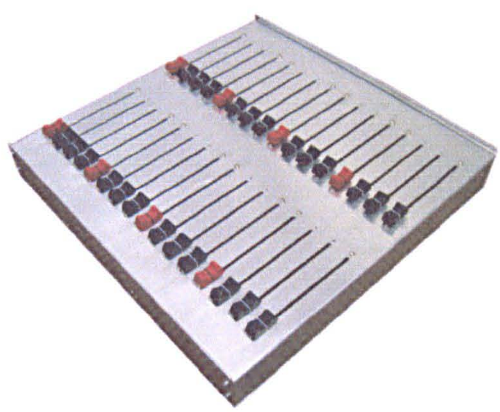

Figure 30. The M2 Sound Diffusion System control surface. Dimensions: 373w x 370d x 55h mm approx.

Audio input and output streams are respectively communicated to and from the system via a MOTU 24 I/O audio interface[239] and this currently determines the number of analogue audio inputs and outputs made available by the system: twenty-four in each case. Again, this device is completely abstracted from the control interface and mix engine and therefore dually serves only as a point-of-entry and point-of-exit for transitorily encoded audio streams. In terms of audio sources, the system includes a Denon DN-C635 professional rack-mounting compact disc player that is permanently

[239] MOTU (2005). "MOTU 24I/O Overview". Website available at: http://www.motu.com/products/pciaudio/24IO/body.html/en.

connected to two of the audio interface analogue inputs. External audio sources can be connected as required to the remaining twenty-two analogue inputs.

The software component of the system is represented by two distinct programs: the *SuperDiffuse Client* application, and the *SuperDiffuse Server* application. The client application effectively allows the performer to specify what each of the thirty-two faders 'does' in terms of diffusion: a number of highly flexible possibilities are available, and these will be described later. Essentially, this information (control data) is communicated in real time to the server application, which performs the function of mixing audio input streams (obtained from the analogue inputs of the audio interface) into audio output streams. The server does this in real time according to data from the control interface received via the client application. Audio output streams are then communicated by the server to the audio interface physical outputs, each of which can be connected directly to an active loudspeaker.

In its present incarnation both client and server applications run on a single computer under the Windows XP operating system. This PC acts, therefore as a central 'communications hub,' insofar as it facilitates communication between all of the discrete components of the system. It is also possible to run the client and server applications on two separate PCs.

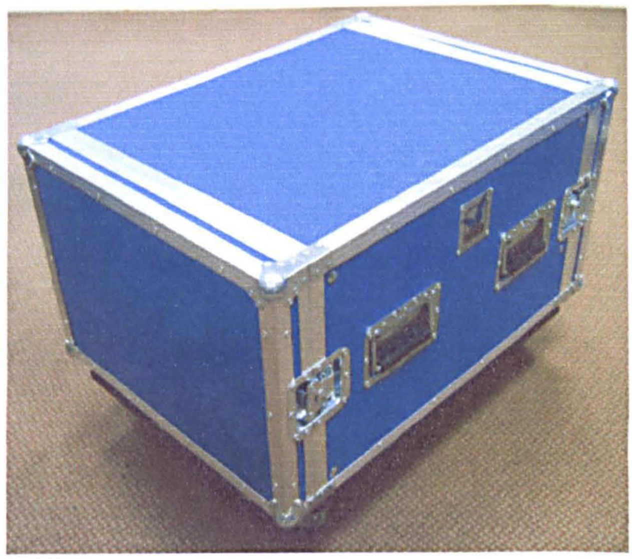

Figure 31. The M2 Sound Diffusion System flight case, containing integrated core system components. Dimensions: 530w x 730d x 540h mm approx.

The core components of the system – this excludes loudspeakers, signal cables, a 15-inch TFT monitor for use with the PC, and any external audio sources required – are flight-cased (see Figure 31) inside a single compact and robust 8U 19-inch rackmount casing. When the detachable front panel of the flight case is removed, the frontal rackmount system – which houses the CD player, audio interface, and 4U rack-mounting PC – can be accessed (see Figure 32). In transit, the detachable rear panel of the flight case houses a wireless keyboard and mouse for use with the rack-mount PC. The rear rackmount system (Figure 33) houses a custom-designed analogue audio I/O breakout panel comprising eight analogue inputs (provided on balanced quarter-inch jack sockets) and twenty-four analogue outputs (balanced XLR connections); these are internally connected to the corresponding inputs and outputs on the audio interface. Access to the back of the PC is also via the rear rackmount system but should not normally be necessary. The rear rackmount system also has a compartment in which the control interface is stored when not in use. A small MIDI interface, wireless keyboard and mouse receiver, and the Ircam AtoMIC are housed internally. Access to these should not be necessary as the design allows for them to remain permanently inter-connected in the appropriate way; however, easy access is available via the front rack-mount system if required. All of the core components obtain mains power via a single standard mains cable, which, in transit, is coiled up inside the flight case and can easily be pulled out via the rear rack-mount system. Cables for connecting the control surface to the Ircam AtoMIC and the monitor to the PC are also provided in this way.

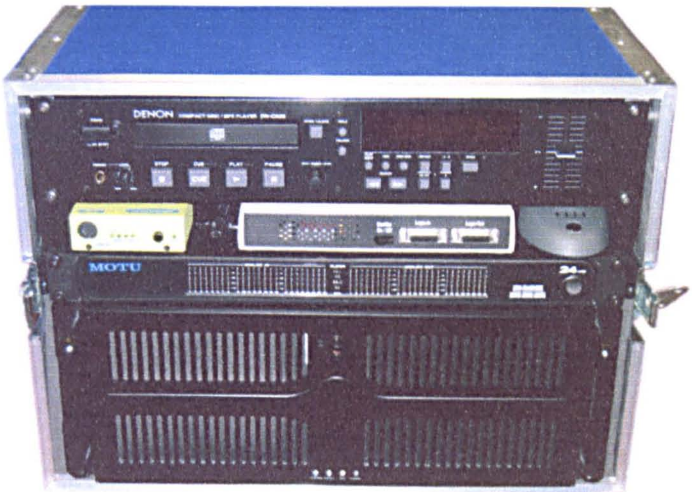

Figure 32. The front rack-mount system, housing (from top-to-bottom/left-to-right): Denon DN-C635 CD player, Midiman USB Midisport 2x2 interface, Ircam AtoMIC Pro v2.0 interface, wireless keyboard and mouse receiver, MOTO 24 I/O audio interface, 4U rack-mounting PC running Windows XP operating system and the *SuperDiffuse* client and server applications.

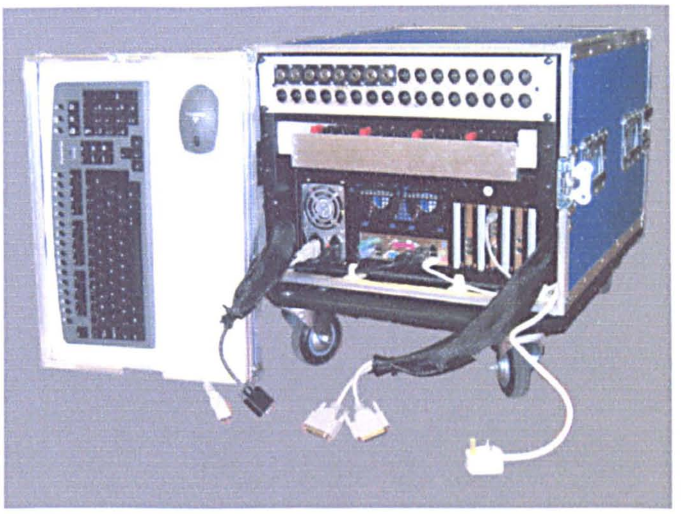

Figure 33. The rear rack-mount system, housing: (from top-to-bottom): custom built analogue audio input/output break-out panel with eight analogue inputs and twenty-four analogue outputs, M2 control surface inside transit compartment, back of rack-mounting PC. Cables shown are (left-to-right) power and video signal for computer monitor, two cables for connection to the control surface, mains power for all of the equipment in the flight case. Wireless keyboard and mouse are also pictured as stored in the detachable rear panel.

5.3. Physical System Architecture

A schematic representation of the M2 Diffusion System, showing the relationship between the core devices, is given in Figure 34, below. Referring also to Figure 20 (page 168) it can be seen that each of the four main diffusion system components (audio source, control interface, mix engine, loudspeaker array) is completely abstracted from the others via various device interfaces. This effectively allows each device to function independently of the others and allows for devices to be interchanged if necessary, without affecting the overall functionality of the system. In other words, the system is modular: we will return to this concept in section 6.2.7. The function of physical audio input layer and physical audio output layer (not labelled in Figure 34) is dually performed here by the audio interface.

It will be noted that the two software components of the system are completely separate entities, communicating via the standard network TCP/IP. This allows for two possible computer configurations, labelled (a) and (b) in Figure 34. In scenario (a) the client and server run on two separate machines. In this case, both machines would obviously require an Ethernet interface to enable communication between them: these are standard on the majority of modern PCs, and in any case can be obtained very cheaply. The client machine would additionally require a USB (or similar, P/S 2 for instance) interface for keyboard and mouse, a VGA socket for connection with a monitor (VDU) – both of these are also standard on all modern PCs at the time of writing – and some kind of MIDI interface for communication with the control surface via the Ircam AtoMIC (the current prototype has a Midiman USB Midisport 2x2 MIDI interface[240]). For the server machine, the only other hardware prerequisite – in addition to the Ethernet connection – is the audio interface itself. The current system, as mentioned previously, uses a MOTU 24 I/O, but any audio interface can be used as long as ASIO compliant drivers are available for it; this has become a widely accepted standard in recent years. No keyboard, mouse, or monitor is required for the server machine as all user interfacing is via the client.

In scenario (b) both the server and client run on a single PC; this configuration is used in the current prototype system. This does not affect the communication between client and server applications, which still happens internally via the TCP/IP. When a single PC is used, it must obviously provide all of the interface devices described. In terms of system functionality, there is absolutely no difference between scenarios (a) and (b).

[240] Midiman no longer exists: the company is now known as M-Audio.

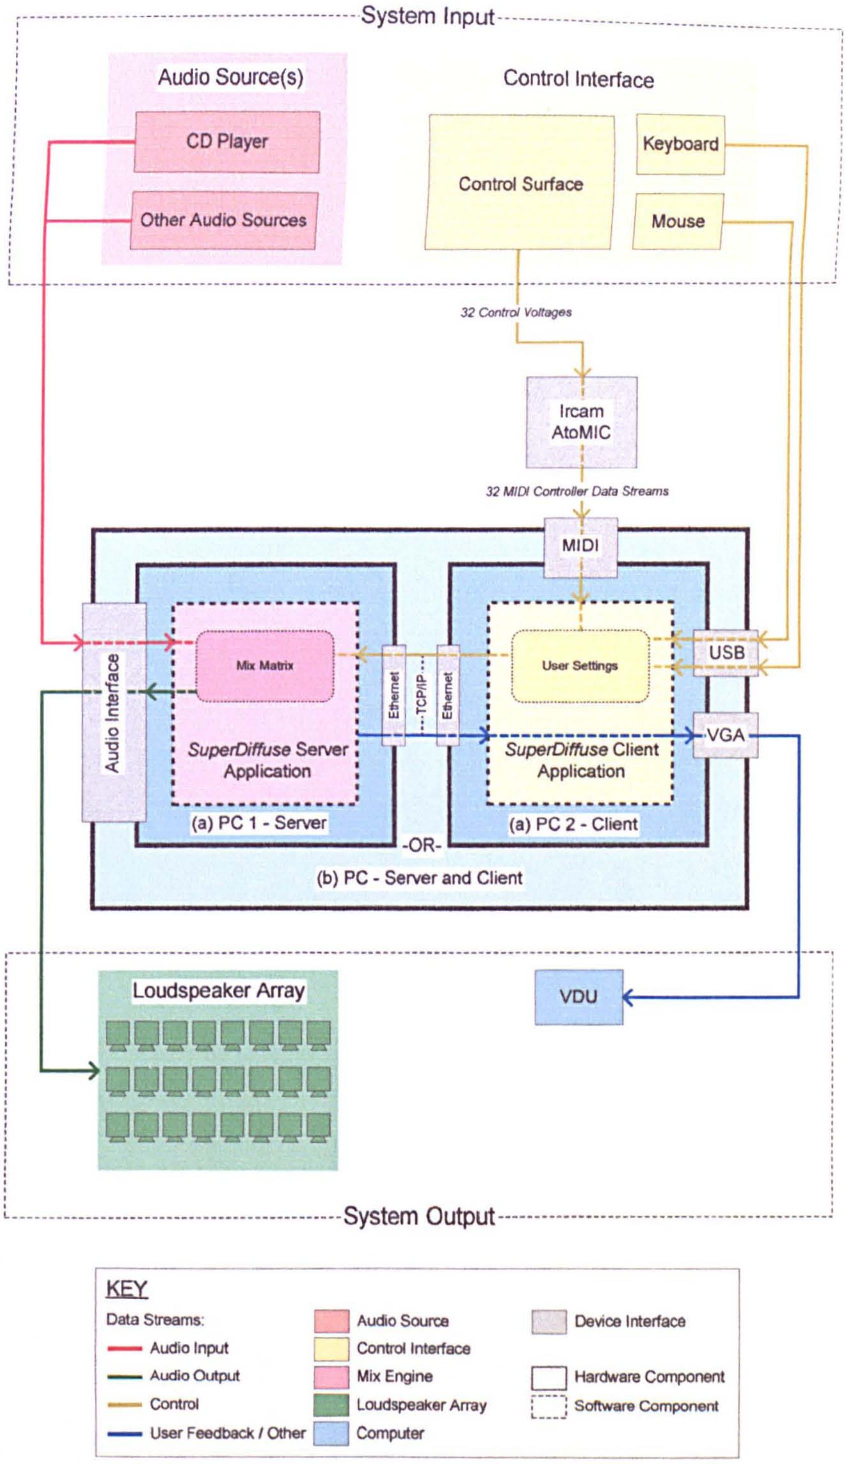

Figure 34. Diagram illustrating the general level architecture of the M2 Sound Diffusion System.

5.4. The *SuperDiffuse* Software

An explanation of the capabilities of the M2 system, in terms of actual sound diffusion, is best approached via an explanation of the capabilities of the software. A comparatively brief explanation of the client and server applications will be given in the following sections. *SuperDiffuse* was designed and coded by Moore, who gives a detailed and comprehensive account of its design and functionality.[241]

5.4.1. *SuperDiffuse* Server Application

In terms of the system as a whole, the *SuperDiffuse* server application is, effectively, the mix engine. This takes the shape of a software model of an audio mixing matrix, as illustrated in Figure 35 (matrix mixers were described, conceptually, in section 4.14). The size of the matrix is determined solely by the number of hardware inputs and outputs available on the sound card and is neither defined nor restricted by the application itself. Each of the matrix crosspoints, inputs, and outputs, has an individual signal attenuator. At present the *SuperDiffuse* server does not support audio 'gain' as such, owing to the risk of digital signal clipping. Audio source signals are routed directly to sound card inputs, and the diffusion data sent from the client application is used by the server to determine the position of each attenuator in real time. In this way, the client and server collaboratively allow any sound card input signal to be dynamically mixed to any sound card output, in any magnitude. In terms of terminology each mix matrix attenuator (inputs, outputs, and cross-points) is considered to be a *parameter* within the *SuperDiffuse* software; there are other parameters that will be discussed in due course.

[241] Moore (2004). "Real-Time Sound Spatialization: Software Design and Implementation". Ph.D. thesis (University of Sheffield).

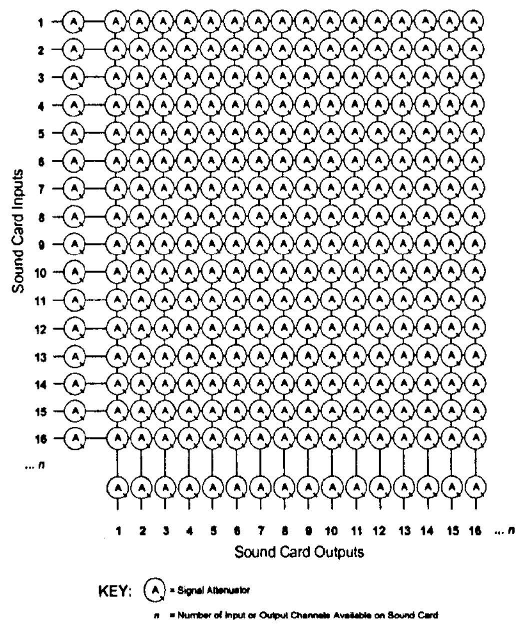

Figure 35. Schematic of the audio mixing matrix emulated by the *SuperDiffuse* server application.

5.4.2. Super Diffuse Client Application

The *SuperDiffuse* client application mediates between the control surface and the *SuperDiffuse* server application and, effectively, represents an abstraction layer between the control interface and mix engine. The 'Performance Window' of the client application (see Figure 36) consists of thirty-two virtual faders, which are directly analogous to the thirty-two physical faders of the control surface. Any movements in the control surface faders will be immediately and exactly reflected on-screen in the Performance Window. Accordingly, the term 'physical fader' can normally be taken mean the to physical fader itself, *and* its on-screen counterpart.

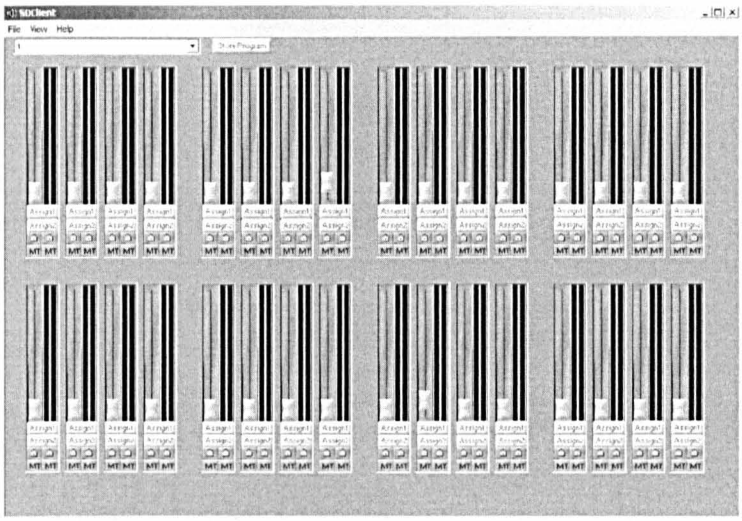

Figure 36. The Performance Window of the *SuperDiffuse* client application.

5.4.2.1. Assigning Mix Matrix Parameters Directly to Physical Faders

In the simplest instance, parameters may be assigned directly to physical faders. This is achieved by clicking one of the 'Assign' buttons in the Performance Window, and then choosing which parameter to assign to the physical fader with the subsequent dialog box (Figure 37). Thus, any one of the physical faders can be used as a direct control for any one of the mix matrix inputs, outputs, or crosspoints. This method alone can be used to emulate the operation of a generic diffusion system. As Figure 36 illustrates, each physical fader has two 'Assign' buttons. Accordingly, each physical fader can in fact have two parameters simultaneously assigned to it. This allows, for example, a single fader to control the output levels of a pair of loudspeakers without the need for an intermediate group fader (these will be described presently). The two independent assignable paths of each physical fader can each be associated with any parameter within the *SuperDiffuse* software, so their flexibility extends beyond simple control of loudspeaker pairs.

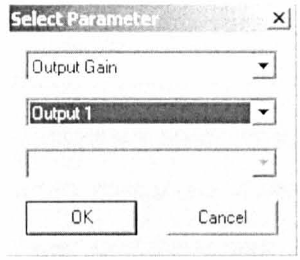

Figure 37. Dialog box with which physical faders are assigned to parameters within the *SuperDiffuse* client application.

5.4.2.2. Group Faders

In addition to the thirty-two physical faders, the *SuperDiffuse* client maintains a bank of 256 virtual 'group' faders. Using the Group Window (the screen-shot shown in Figure 38 shows two audio inputs and six audio outputs, reflecting the profile of the audio interface in use at the time) each of these virtual faders can be configured to define the respective values of a collection of parameters. To this end, the Group Window displays a diagram of the mix matrix in which each element can be associated with a value ranging from -127 to 127 inclusive.

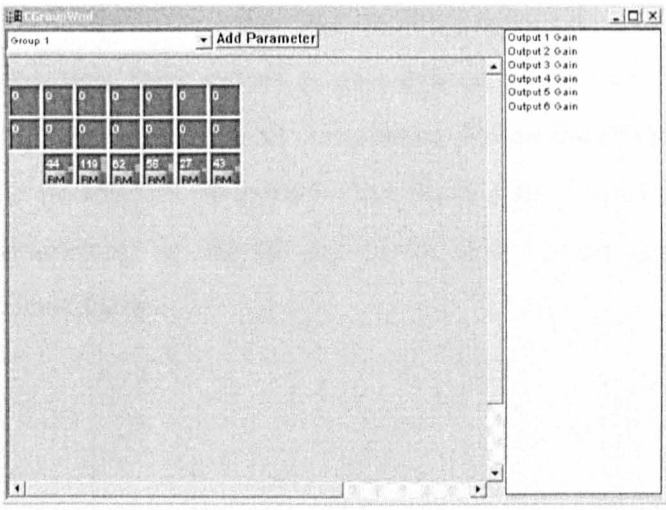

Figure 38. The Group Window of the Super Diffuse Client application.

Once defined in this way, the value of each virtual group fader becomes a parameter within the software, and can therefore be associated with a physical fader assign path. When this is done, the parameter values contained within that group fader will be reflected when the corresponding physical fader is at the 'maximum' position, with intermediate physical fader positions resulting in grouped parameter values being scaled linearly. Note that, within groups, parameters can be also be given negative values.

5.4.2.3. Effect Faders and Effect Types

The *SuperDiffuse* client also maintains a bank of 256 virtual effect faders, which can be configured via the Effects Window. Using this interface, an effect fader is associated with an effect type. At present there are three different effect types – Wave (pictured in Figure 39), Chase (Figure 40), and Randomise (Figure 41) – each with its own set of user-configurable characteristics. Each effect type has its own unique behaviour, directly affecting two or more parameters.

The Wave Effect takes one or two parameters, and sinusoidally modulates their values at user-defined frequencies, phases, and amplitudes. As with all assignments within the client application, the parameters assigned to the Wave Effect can be mix-matrix parameters, or indeed any of the 256 virtual group or other effects faders.

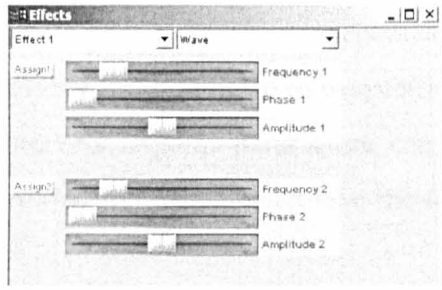

Figure 39. The Effects Window of the *SuperDiffuse* client application, showing a Wave Effect assigned to virtual Effect Fader 1.

The Chase Effect is a basic twenty-four-step sequencer and, as such, can be given up to twenty-four parameters. Each parameter is in turn associated with a value between 0 and 127 (inclusive), and the Chase Effect steps through the sequence of parameter-value settings at a user-defined rate, with cross-fades automatically implemented between steps. This effect can be used to sequence mix-matrix parameters and, more interestingly, group and effect faders, or any combination thereof.

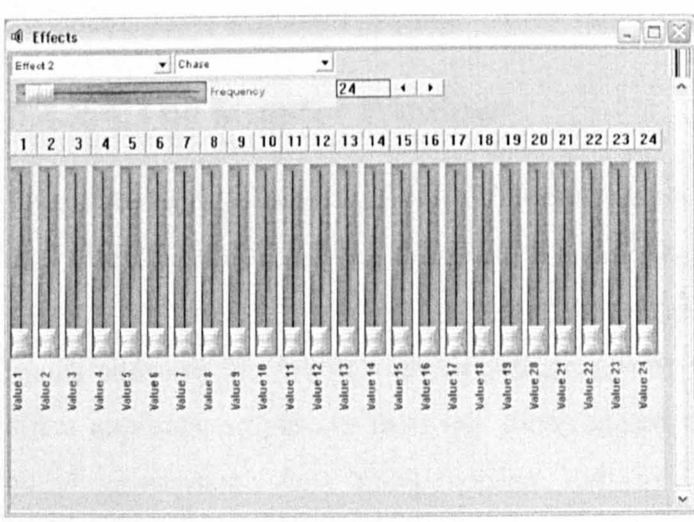

Figure 40. The Effects Window of the SuperDiffuse client application, showing a Chase Effect assigned to virtual Effect Fader 2.

The Randomise Effect takes one or two parameters and randomises their values independently – additively, subtractively, or both – within user-specified ranges and at user-defined intervals. Again, parameters can be obtained directly from the mix-matrix, or they can take the form of group or effect faders.

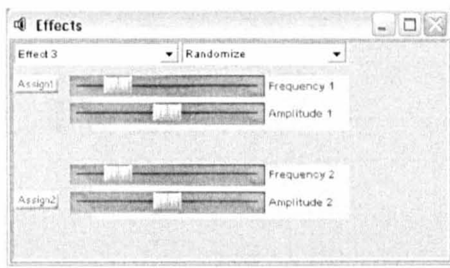

Figure 41. The Effects Window of the *SuperDiffuse* client application, showing a Randomise Effect assigned to virtual Effect Fader 3.

Once an Effect Fader has been associated with an Effect Type, the value of the Effect Fader becomes a parameter, which can then be assigned to a Physical Fader, or to any of the other assignable paths within the system.

5.4.2.4. The Monitor Window

Levels for each parameter of the mix matrix can be observed at any time by means of the Monitor Window. Note that the Monitor Window does not provide signal levels for sound card inputs and outputs: this is a physical impossibility because the client application does not have any direct means of accessing audio processing, this being carried out by the server. Additionally, mix-matrix parameters can be muted or 'locked' from the Monitor Window. The value of a muted parameter will always be zero, while the value of a 'locked' parameter will always be maximum (127), regardless of any assignments made to that parameter. Some reasons why this might be necessary will be described later, in section 5.5.3. Suffice, at this point, to say that the ability to lock parameters of the mix matrix allows for sound diffusion to be executed in a number of different ways by allowing the user to define the routing limitations (of course, there need not be any at all) in a manner that will be advantageous to the type of diffusion required.

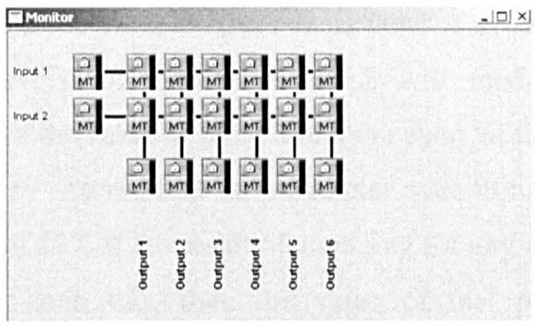

Figure 42. The Monitor Window of the Super Diffuse Client application. In this case, a sound card with two audio inputs and six audio outputs is being used.

5.4.2.5. Summing of Parameters

It should be clear that, owing to the flexible nature of parameter assignments within the system, it is possible that any given parameter could conceivably have its value modified from several 'places' simultaneously; indeed this is common practice in the operation of the M2 system. Consider the following scenario. Matrix inputs 1 and 2, and crosspoints 1-1 and 2-2 are locked. Matrix outputs 1 and 2 are assigned to a Physical Fader *A*. Matrix outputs 1 and 2 are *also* assigned as the two parameters of a Wave Effect, associated with one of the 256 virtual Effect Faders. This Effect Fader is, in turn, assigned to Physical Fader *B*. On the Control Surface, we now have two Physical Faders that are able to exert control over matrix outputs 1 and 2: Physical Fader *A* can be used simply to control the level of inputs 1 and 2 that gets sent to outputs 1 and 2; Physical Fader

B will modulate the 'levels' of matrix outputs 1 and 2 sinusoidally.

So, what happens if we raise the levels of both Physical Faders *A* and *B* on the Control Surface? The target parameter values are summed. If we raise only fader *B* (and fader *A* is at 'zero'), the 'levels' of outputs 1 and 2 will modulate sinusoidally between zero and the value of fader *B*. If fader *A* is already raised, then the 'levels' of outputs 1 and 2 will modulate sinusoidally between the value of fader *A* and the value of fader *A* plus *B*. The software ensures that no parameter ever exceeds its maximum value of 127: if the result of summing for any given parameter is greater than this, then the value of that parameter will be 'clipped' at 127. This does not mean that the audio signal clips (in fact, this safety feature is included to ensure that the audio signal does *not* clip), but merely that the parameter will not exceed its maximum value.

5.4.2.6. Summary

To summarise, the *SuperDiffuse* client application presents a number of paradigms by which the server's mix matrix can be controlled via the M2 Control Surface interface. Mix matrix parameters may be controlled on a simple one-to-one, or two-to-one, basis with Physical Faders. Alternatively, virtual Group Faders allow multiple parameters to be grouped and controlled proportionally, while Effect Faders allow groups of parameters to be automated in various different ways. Generally speaking, any parameter can be assigned to any entity that requires parameters, as illustrated in Figure 43. The exception to this is with Effect Types; these may *only* be assigned to Effect Faders, at which point the Effect becomes the (only) means for the Effect Fader to accept parameters.

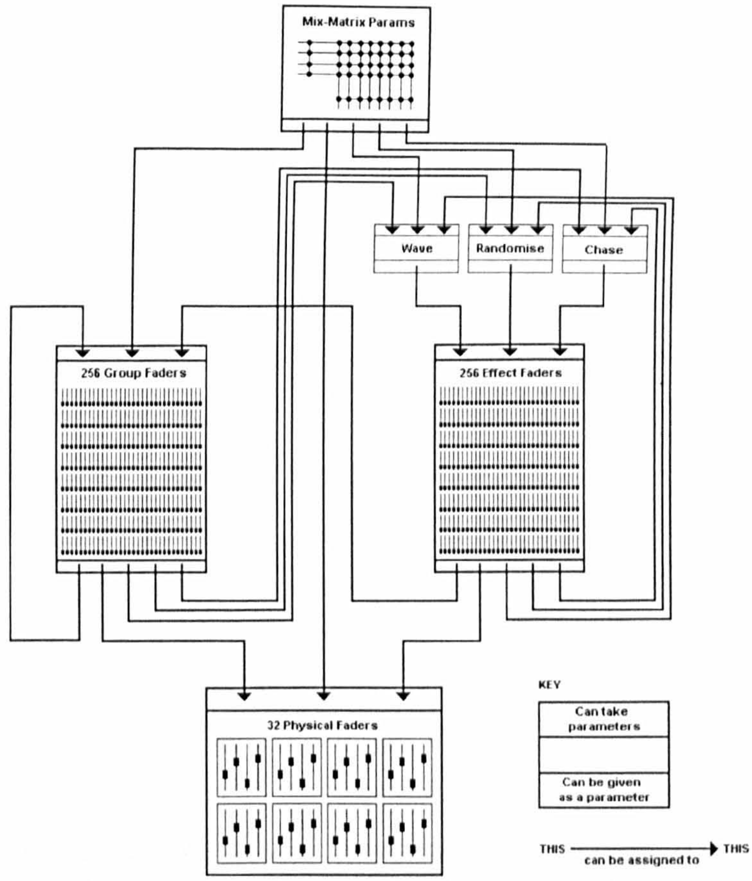

Figure 43. Summary of assignable paths for parameters within the *SuperDiffuse* client application.

5.5. Using the System

The following sections describe the typical process involved prior to staging a performance of electroacoustic music using the M2 Sound Diffusion System.

5.5.1. Transportation

The M2 has been specifically designed for easy transportation to venues. The flight case measures 53cm wide by 73cm deep by 54cm tall and can fit into the back of most cars if this is necessary. The flight case is heavy but has wheels and so can very easily be moved on the ground by one person, but two people are realistically needed to get the unit into and out of a vehicle, or if stairs are involved. The wheels of the flight case can lock for safety during transit. Aside from the flight case, it will also be necessary to transport loudspeakers, stands, at least as many signal cables and standard IEC mains connectors as there are loudspeakers, a number of mains power splitters to power the loudspeakers, and a second very small flight case containing a flat-screen TFT monitor. If only diffusion from compact disc is required (this is at least relatively often the case), this is all that is required. Reference to Figure 34 should indicate what other components will be needed for larger performances; it should be noted that the current prototype does *not* include microphone pre-amplifiers.

It will be clear that, in the vast majority of cases, the bulk of the equipment for transportation will be represented by loudspeakers (in the author's experience this is also the case with most commonly employed diffusion systems), and purely on this basis a van (or several car journeys) is often required except where small loudspeaker arrays are used. There are also likely to be a number of small sundry items such as desk lamps and so on.

5.5.2. Physical Setup

The flight case is normally the first piece of equipment to be installed, and is placed close to the desired location of the diffusion console. In performance all of the core system components remain housed inside the flight case, and on several occasions the flight case itself has been used as a table for the control surface. Although amply large enough for the control interface alone, this is not, however, ideal as it is difficult to accommodate the computer monitor, despite that fact that this is flat-screen, and on this basis a table is normally required.

When in place, the front and back detachable panels of the flight case are removed. The keyboard and mouse are removed from inside the rear detachable panel, and the batteries for each installed. The front detachable panel has been used in the past as a stand for the control surface, neatly providing a raked angle towards the performer if this is preferred. The control surface slides out from its compartment inside the flight case. The mains cable for the flight case is uncoiled via the open back panel. As described previously, the flight case components are all powered from a single mains connection, and so only one mains socket is required. Cables to connect the control surface and computer monitor are also uncoiled from inside the flight case and connected. The small flight case used to house the monitor in transit doubles as a stand for the monitor. The keyboard and mouse are wireless, so it is not necessary to physically connect these.

The placement of loudspeakers on stands throughout the venue normally takes place next, followed by the provision of mains power for each of these. The time taken to achieve this invariably outweighs the set up time for the flight-cased components by a considerable margin. With the loudspeakers in place, each of these is connected via a standard balanced XLR signal cable to the audio interface via the break-out panel in the rear rack-mount system of the flight case, a note being made of which loudspeakers are connected to which audio interface outputs. When this is completed, the audio interface is switched on, computer booted up, and the *SuperDiffuse* server and client applications launched in that order. Loudspeakers are then switched on, completing the physical setup of the system.

5.5.3. Software Configuration

With the system in place, the software is then configured according to the requirements of the performance. If a preset has been saved from a performance using an identical audio interface and loudspeaker array, this is simply a matter of loading the corresponding file. If not, the system must be configured.

The *SuperDiffuse* server application does not require any configuration as this is all done via the client. In configuring the client application we are ultimately determining the way in which the control surface can be used to execute diffusion and, therefore, this stage of the setup process is very important. Two procedures are common: either the same configuration is used for every piece in the concert programme, or else different pieces have different configurations. The former case is only really possible if all of the pieces have the same technological profile in terms of physical audio sources and CASS specifications. This scenario would be possible, for example, if all of the pieces were to be diffused from stereophonic CD. For the purposes of explanation it will be assumed that we are dealing with this scenario. In any case, where different configurations are required, the same process has to be iterated again for each new configuration.

Firstly, a decision must be made with regard to how the signal routing within the system will operate. In terms of the mix matrix (see Figure 35, above), here we are effectively deciding which parameters will be controlled via the control surface, and which (if any) are to be locked to maximum amplitude. Two different routing examples for the pluriphonic diffusion of stereophonic works will be given, as this is conceptually simple. The same basic principles would, of course, apply for larger CASSs and different kinds of diffusion. The examples given will assume, again for simplicity, that the BEAST 'main eight' array is being used, and that audio interface outputs are connected to loudspeakers as represented in Figure 44.

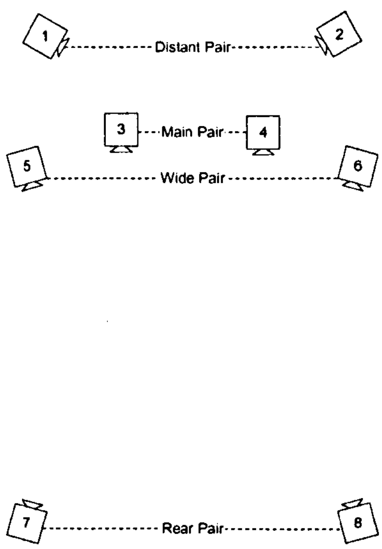

Figure 44. Loudspeaker array and audio-output-to-loudspeaker connections assumed for the two *SuperDiffuse* routing configuration examples given below. A stereophonic audio source from CD is also assumed.

In the first example, matrix inputs 1 and 2, and cross-points 1-1, 2-2, 1-3, 2-4, 1-5, 2-6, 1-7, and 2-8 are locked, as shown in Figure 43, below. This means that audio inputs 1 and 2 are effectively 'hard wired' to audio outputs 1 and 2, 3 and 4, 5 and 6, and 7 and 8, respectively. In order to achieve diffusion to loudspeakers, therefore, audio outputs 1 to 8 will have to be assigned to physical control surface faders. In the simplest instance, this could be done by assigning each of the eight mix matrix output parameters to an individual fader. The fader to which output parameter 1 is assigned, would afford control over the output level of the left-distant loudspeaker, referring to Figure 44, the fader controlling matrix output 2 would afford control over the right-distant loudspeaker, and so on. The kind of diffusion control made available by this configuration (one-fader-to-one-loudspeaker) will be familiar to most practitioners.

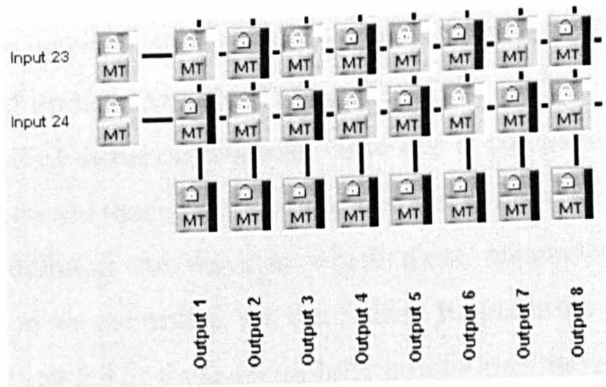

Figure 45. First routing example. Matrix inputs 1 and 2, and cross-points 1-1, 2-2-, 1-3, 2-4, 1-5, 2-6, 1-7, and 2-8 are locked, leaving matrix outputs 1 to 8 available for assignment to physical faders.

In the second example, matrix inputs 1 and 2, and matrix outputs 1 to 8 are locked, as illustrated in Figure 44, below. In this case, no audio inputs are 'hard wired' to outputs as such, and in order to hear anything we will need to be able to control the mix matrix cross-points. The same kind of control attained in the first example could be achieved in this case by assigning cross-points 1-1, 2-2, 1-3, 2-4, 1-5, 2-6, 1-7, and 2-8 each to an individual fader.

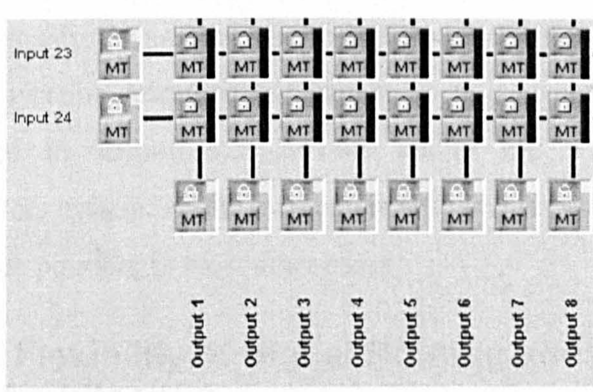

Figure 46. Second routing example. inputs 1 and 2, and matrix outputs 1 to 8 are locked, leaving matrix cross-points 1-1, 2-2-, 1-3, 2-4, 1-5, 2-6, 1-7, and 2-8 available for assignment to physical faders.

Clearly, the nature of the routing configuration in the Monitor Window will depend on the technological profile of the work to be performed, and the kind of diffusion required. Once this is completed (and, indeed, it may be decided that no mix matrix parameters are to be locked), it is a matter of defining the ways in which those parameters that are not locked are to be controlled via the faders. It should be clear, referring back to section 5.4.2, there are infinite possibilities ranging from direct one-to-one control of loudspeakers via faders (as described in the previous examples), through to configurations where the relationship between faders and mix matrix parameters is highly abstracted.

Once the software has been set up, the configuration can be saved as a file. A *SuperDiffuse* performance file can contain up to 128 preset programs, each of which can be configured differently via the process accounted above. During performance, programs can be selected from a drop-down menu (this can even be done while the server is processing audio, without any glitches). This method would be adopted in cases where different configurations are required from one piece to the next.

5.6. Evaluation

As stated previously, evaluation of the M2 Sound Diffusion System – as will all of the systems described in Chapter 4 – will be according to those criteria defined in section 4.3. As the author has a comprehensive knowledge of this system, so the evaluation in this case will be systematic: this has not been possible in most other cases.

5.6.1. Flexibility of Signal Routing and Mixing Architecture

At its heart, the M2 Sound Diffusion System operates on a matrix-based mixing architecture. As explained in the previous chapter, this is the most flexible mixing architecture available at the present time, because it allows any input to be mixed to any output in any quantity. It is also, however, reasonably convoluted, potentially confusing, and extremely difficult to control directly in the context of sound diffusion, owing to the number of parameters and the fact that these do not always correspond in obvious ways to the desired outcomes. In short, the matrix mixing paradigm has the potential to be enormously non-ergonomic. The extent to which this has been successfully addressed in the M2 system will be evaluated in the next section.

In terms of signal routing, the M2 is certainly among the most flexible of the systems discussed, because this can be defined by the user, in software. In this way, well-established paradigms – such as bussing in stereo pairs – can easily be configured, but the system is not limited to these. The client application also allows for signal routing presets to be stored and recalled instantly, meaning that each item within a concert programme can use an appropriate routing. This facility is not available in any other system as far as the author is aware, and therefore represents one of the considerable strengths of the M2.

This means that physical re-patching during a performance should never be necessary, so long as the total number of audio input channels used throughout the programme does not exceed twenty-four (of which two channels must be from CD). Put simply, it does not matter which audio inputs the audio sources are connected to, because the routing is all determined in software. However, the use of more than eight non-CD audio input channels – although perfectly possible – is not as easy as it could be, owing to the fact that only eight inputs are externally accessible on the rear rack-mount system; this can increase set-up time slightly for concert programmes requiring more than eight non-CD inputs. In addition, any audio sources that cannot be routed to standard XLR connectors will require adapters, although this criticism can probably be filed against any current diffusion system.

The total number of audio outputs presently offered by the system – twenty-four – allows for a reasonably generous loudspeaker array to be used in concerts. The system has been used to its full capacity in this respect in at least three public performances and has performed flawlessly. A twenty-four channel array, however, is not by any means large in comparison with several of the systems described previously, and does not rival the Acousmonium, Cybernéphone, BEAST, or sound cupola systems.

The number of audio inputs and outputs is, however, flexible, because further modules can be added to the MOTU PCI-424 PCI card. At present, further 24 analogue I/O modules are available from MOTU, as well as the MOTU 2408 Mark 3, which adds eight analogue inputs and outputs and twenty-four digital ins and outs via A-DAT optical and Tascam T-DIF interfaces. A maximum of three further modules can be added, allowing a number of possible permutations. The maximum number of input and output channels attainable with this technology is ninety-six of each, with various combinations of analogue and digital possible. This exceeds the number of inputs and outputs offered by all but one of the systems described in the previous chapter, the exception being certain of the sound cupolas (see section 4.12). In addition, the system will function correctly with any sound card that can be accommodated by the PC's motherboard (normally these will have the standard PCI interface) provided that the sound card in question has ASIO drivers, as many PCI sound cards nowadays do. This will prove to be advantageous in a number of ways to be explained later.

Another advantage is that the control interface is abstracted from the signal routing and mixing architecture. Therefore, the design of the architecture itself has not been limited in any way by control considerations. This results in a higher degree of flexibility in the present context, and is also beneficial in relation to interface ergonomics.

5.6.2. Interface Ergonomics

On the surface, the M2 system would appear to be based around a rather unoriginal bank of thirty-two standard faders. Issues relating to this, in terms of interface ergonomics, have been described previously.

However, most of these issues arise when the faders are themselves an integral part of the signal routing and mixing architecture, and this is not the case with the present system. This means that *more ergonomic use* (or more 'efficient' use) can be made of the faders, which also have the considerable advantage of being an interface familiar to many practitioners (this latter point will be discussed more fully in the next section).

Among the many apparently straight-forward advantages offered by this design is – quite simply – the ability to control the loudness of a *pair* (or, indeed, a group of arbitrary denomination) of loudspeakers using only a single fader. This can be achieved by assigning the appropriate mix-matrix parameters to a virtual group fader, which is, in turn, assigned to a single physical fader. With some of the more widespread systems (including both the generic systems and the BEAST system) it is necessary to interact with two faders (one for each loudspeaker) to obtain this result. In cases where both faders will be manipulated identically – and in many cases preserving the (e.g.) stereo balance in this way will indeed be a concern – the ability to carry out this action with a single fader most certainly represents an ergonomic improvement.

A unique feature of the M2 system – along the same lines – is the possibility of realising proportional amplitude control over a group of loudspeakers; neither the Acousmonium nor the BEAST experimental system offer this possibility. In this way, raising a fader to its 'maximum' position could – for instance – simultaneously bring one pair of loudspeakers to maximum output level, a further pair to 50% maximum level, a further pair to 30%, and so on. Of course, the exact nature of the proportions is completely arbitrary (and loudspeakers need not be grouped in pairs if this is not appropriate), and far more abstract groups (comprising further groups and effect faders in addition to direct mix-matrix parameters, for example) can be defined if required. Assigning grouped parameters with negative values can yield interesting results. Let us imagine that the system is in a state whereby every loudspeaker in the array is raised to maximum output level. If a virtual group fader (assigned to a physical fader) has each of the matrix output channels assigned to it with a negative maximum value, raising that fader to its maximum position will *lower* the output levels of all the loudspeakers to zero. This technique has been implemented in many performances with the M2, a good example being the use of a single fader that brings the diffusion 'all to mains' (that is, only a frontal pair of loudspeakers remain active). Such techniques are only available owing to the fact that mix-matrix parameter values are summed, as described previously, and this feature is not available on any widely-used diffusion system. In all of the examples given, to achieve the same result by manipulating one fader per loudspeaker would be considerably less ergonomic, and in many cases infinitely less ergonomic (i.e. impossible).

Overall, the nature of the assignments that can be made, and the effects these have on the behaviour of the control interface with respect to the diffusion, is extremely open-ended, and the user has direct control over this. To this end, the user can tailor the controls to be ergonomic with respect to the specific diffusion actions required. This would appear to be a rare commodity indeed, with most existing systems offering only a fixed repertoire of techniques, and is certainly one of the strongest benefits of the M2 system in the present context. Additionally, the configuration can be instantaneously changed from one piece to the next (or, indeed, at any time) by recalling a previously defined preset.

An exploration of the ways in which the use of faders has been ergonomically improved aside, it could be argued that because the system adopts the fader as its only interfacing paradigm (the keyboard and mouse are only used to configure the system and not for actual diffusion), it is therefore – on a physical level – still subject to the ergonomic limitations thereof, progressive control over one dimension only being a good example. To put it another way, while the M2 may offer extended functionality, the physical mode of operation remains the same. As described in Chapter 1, every creative framework has its own particular characteristics and limitations, and the degree of restriction imposed by the interface as a whole could be lessened with the inclusion of alternative paradigms (joysticks, buttons, foot-pedals, etc.) alongside more conventional means. A fuller exploration of such possibilities would certainly be beneficial, and is not evidenced in the system as it stands.

The most obvious ergonomic limitations of the M2, however, reside in the way in which the system itself must be configured by the user, via the client application's graphical user interface (GUI). The way in which assignments are made, for instance, is relatively inefficient in comparison with what it could be, and many optimisations could be made in this respect. As a representative example, consider the process of configuring a physical fader to control the output levels of a pair of loudspeakers in equal proportion. This procedure would involve (depending, of course, on the configuration of the Monitor Window – see Figure 42; one possibility will be described here) clicking on the physical fader's 'Assign' button, selecting 'Output Level' from a drop-down menu, selecting *which* output level (i.e. which mix-matrix output attenuator) from a further drop-down menu, and then clicking 'OK' to confirm the assignment. This process would then have to be repeated for the physical fader's second 'Assign' button (recalling that we are assigning two output levels to a single control), resulting in a total of twelve mouse-clicks to complete the assignment. For the most part, all assignments must be made in this way, and the example given is a relatively simple one. With more complex scenarios involving the assignment of parameters to multiple virtual group and effect faders *prior* to physical fader assignment, it is clear that configuring the system has the potential to become a tedious and time-consuming process. This is not owing to the nature of the task being complex, as such, and therefore represents poor interface ergonomics. With available rehearsal time usually limited to begin with, this could be a serious problem.

It can therefore be seen that, in making enormous ergonomic improvements in certain areas, the M2 system consequently poses a different set of ergonomic challenges in other areas, and some of these are less well dealt with. Nonetheless, the problems are, for the most part, in relation to configuring the system prior to performance, with ergonomics *during* performance substantially improved in comparison with other systems.

5.6.3. Accessibility with respect to Operational Paradigms

Although cited – in one sense – as a potential limitation in the previous section, the M2's use of faders represents a considerable advantage under the present heading. As described in the previous chapter, familiarity is the single greatest aid to the realisation of an accessible interface, and the mixing desk fader is most certainly familiar in the context of sound diffusion. This was, for this very reason, an explicit choice in the design of the control surface, and 100-millimetre throw ALPS faders – frequently used in commercially available audio mixing desks – were specifically selected to enhance the sense of familiarity, despite the fact that the control surface does not have any direct contact with encoded audio streams. This particular design choice was a good one, feedback from users frequently indicating that the interface does not 'feel' substantially different from a conventional mixing desk.[242] In conjunction with the increased functionality of this already familiar interface (it is possible to 'do more' with the faders than in certain other systems employing the same physical interface), this approach has been very accurately described by Lewis as 'supplementing existing diffusion techniques rather than replacing them,'[243] and this has proven to be useful in the present context. From a familiar starting-point, practitioners are able to progressively – over the course of several

[242] Although Robert Normandeau is reported to have expressed a distaste for these faders, on the grounds that they do not offer enough resistance. Nikos Stavropoulos (2005). Personal communication with the author.
[243] Dr. Andrew Lewis, University of Wales, Bangor (2004). Personal email communication with the author.

performances and rehearsals, say – extend their repertoire of diffusion actions through more and more complex assignments to the physical faders. This is not true of other systems, which either tend to provide the performer with a relatively limited selection of potential techniques and fixed functionalities, or with paradigms that are less familiar, or both.

The visual and tactile familiarity of the interface has also, however, proved problematic in other respects, certain users – for instance – finding it difficult to comprehend the fundamental difference between the M2 control surface and a conventional audio mixing desk, as implemented in a generic diffusion system. This is due at least partly to the fact that the system *can* be configured (and quite often is) to behave in this way (faders eliciting one-to-one control over loudspeakers, *et cetera*), and when this is the case, the results are highly convincing. When presented with an interface that looks, feels, and behaves in more or less the same way as a conventional mixing desk, it is hardly surprising that practitioners get confused from time to time.

This problem is also due, albeit more indirectly, to certain of the ergonomic inefficiencies described in the previous section. Because the process of configuring the software is somewhat cumbersome, and because there is usually little time available to train all of the performers in the necessary procedures (this raises further questions with respect to accessibility, which will be discussed more fully in the next section), so many concerts have, in the past, been rehearsed and performed with the system pre-configured by an expert user. This does, in certain respects, defeat the object of having such a highly configurable system, and has fairly often resulted in practitioners taking the 'easy option' and using the M2 as if it were a generic system, effectively bypassing many of its advantageous features.

This circumstance is both unbeneficial and self-perpetuating insofar as practitioners are only relatively rarely able to gain first-hand experience of the more configurable aspects of the system, and therefore have little opportunity to assimilate the skills necessary to take advantage of these for their own individual purposes. It is for this reason that the abstract, open-ended, nature of the system is often remains ill understood by new users, which is unfortunate because this is, without a doubt, one of the system's strongest features.

For these reasons – apparent similarly to conventional mixing hardware, and a lack of knowledge with respect to the fundamental concepts and potential capabilities of the system – it has been noted that certain users become disorientated because the system does not behave as they have expected. When, for instance, raising a fader results in the *reduction* of loudspeaker output levels, and it is not immediately apparent to the performer why this has been the case, the tendency – particularly during performance – might be to panic. From this perspective, the behaviour of the system can appear highly erratic and unpredictable, and it is therefore not surprising that certain practitioners opt out of the more unconventional capabilities. This does make it somewhat difficult to engender a culture of 'supplementing existing diffusion techniques rather than replacing them,' and in this respect it could be argued that 'throwing performers in at the deep-end' with completely unfamiliar paradigms – and without a safe, familiar, option – might be more successful.

Overall, however, the M2 exhibits the potential for a high degree of accessibility with respect to the nature of its operational paradigms, and positive feedback has been received on this basis. Some suggestions as to how problems noted in the present context can be indirectly addressed will be described in the next section.

5.6.4. Accessibility with respect to Technological System Prerequisites and Support thereof

The constituent parts of the M2 system can be observed in Figure 34 (page 247). This summarises the software and hardware components required to build the current M2 prototype, excluding flight casing. It can be seen that many of the components are very easily, readily, and comparatively cheaply available. These include the CD player, audio interface, computer and computer-related hardware components, and the loudspeaker array (which in any case would cost the same for any sound diffusion system). It will also be noted, however, that the system comprises certain bespoke components, including both software applications and the control surface. It also incorporates an Ircam AtoMIC interface, which, at a cost of around US$695,[244] is excessively expensive in this context, given that its sole function is to perform control voltage to MIDI conversion. This means that the current M2 prototype is not as accessible as might be desirable.

However, the requirements described represent only one potential configuration. The M2 deliberately offers a high degree of modularity, and this is strongly beneficial from the point of view of accessibility in the present context. All of the devices are completely abstracted from each other, and communication between them achieved via standardised protocols, effectively meaning that every component of the system (realistically excluding the software, which is essential to maintaining the core system functionality; this will be discussed presently) is interchangeable, so long as the appropriate communication protocols are observed. This means that the system can be tailored in accordance with a number of important criteria – including cost, potential application, and personal preference – without affecting the fundamental functionalities of the system as a whole. The expensive Ircam AtoMIC can be dispensed with, for example, if an alternative MIDI device is used to control the diffusion. Indeed, before the development of the control surface, the system was extensively tested using a Kenton Control Freak 16-fader MIDI interface, and many other alternative devices are available commercially. The same interchangeability can be

[244] Retail price in 2000, as cited in:
Cutler, Robair and Bean (2000). "The Outer Limits - A Survey of Unconventional Musical Input Devices". *Electronic Musician*. Electronic journal available at: http://emusician.com/mag/emusic_outer_limits/.

observed in many of the other system components, and this has been a conscious design feature.

Another important advantage rests in the fact that the system is highly scalable. Owing to its modularity, the system can be scaled down (or up) very effectively, and this is highly beneficial insofar as a scaled down version of the system can effectively be 'taken home' for practice purposes. However, at present – owing in part to the nature of the software interface – system operation does depend considerably on the loudspeaker array being used, in terms of number of loudspeakers and their relative positioning. In this respect, much could be learned from, for instance, the DM-8 and ACAT Dome systems, which, to a certain extent, allow arrays of various sizes to be dealt with in essentially the same way.

Use of the M2 for rehearsal outside the performance context is also hindered by a lack of access to the software components, which – at the time of writing – are neither available commercially nor for free. At present this precludes the use of the M2 system in a wider context and this is a problem that urgently needs to be addressed. Some possible solutions will be proposed in the next chapter. Overall, however, although the components of the M2 (the software in particular) may not be as immediately accessible as those required to build, say, a generic diffusion system, the functionalities offered are considerably extended, and the system has the notable advantage of being highly scalable without unduly affecting this.

5.6.5. Compactness / Portability and Integration of Components

A great deal of attention has been paid to the efficient integration of components within the M2 system. Because many of the core elements can remain permanently inter-connected and housed within the flight case, the system can be transported with comparative ease and set up very quickly. Once the flight case is in place in the venue, the core of the system can be physically set up by one person and fully functioning in well under five minutes,[245] provided the integrated CD player is the only audio source required. This does not take into account, however, the time taken to deploy and power up loudspeakers and connect these via signal cables to the audio interface outputs, which – it can be argued – represents a considerable portion of the setup overhead for any portable diffusion system. If any audio sources other than compact disc are required, then the appropriate hardware must be integrated, and it has also been noted that software configuration of the system can be time-consuming.

The portability of the system is also greatly improved by its scalability, as described previously. While there are doubtlessly further optimisations that could be made to decrease setup time, in comparison with other systems the M2 performs extremely well in this respect, and – importantly – the portability of the system does not compromise its overall flexibility.

5.6.6. Reliability

In having hosted over a dozen public concerts (at the time of writing), and numerous live demonstrations and workshops since its inaugural performance in March 2004, the software component of the M2 system has not crashed once during performance or rehearsal. On one occasion the system was left set up and running continuously for around sixty hours, in which time three full days of rehearsal and three public evening performances took place, involving multiple performers.[246] In this case, all twenty-four output channels were in use, and no problems were experienced.

[245] This assumes that the user is familiar with the procedure. The system has been set up in tests at the University of Sheffield Sound Studios – from fully flight-cased to up-and-running – in just over two minutes.

[246] USSS Sound Junction III was staged at the University of Sheffield Drama Studio between Thursday 24th and Saturday 26th February 2005, and featured – amongst others – composers Francis Dhomont and Andrew Lewis performing a selection of their own works using the M2 system.

If, at any point, TCP/IP communication between the client and server applications is lost, the server application will maintain the state of the mix-matrix and this will therefore not result in the loss of sound projection, but only a loss of control over the mix matrix. The work in question can therefore continue to play until the end, at which point the client application can be restarted. To date, this has never happened, except deliberately during testing. Testing has also proved that if communication with the client is lost, the client can be restarted, and communication with the server re-established without any disruption to the DSP engine; diffusion can then be continued as normal. Various other tests have also been carried out in which – for example – multiple copies of the client application have been loaded, software settings altered, saved, and re-loaded, complex signal routing configurations defined in real time, and so on, and none of these has resulted in any disruption whatsoever to the audio playback.

However, the software has not been tested on as wide a variety of different hardware configurations (different audio interfaces, different PC software and hardware configurations, *et cetera*) as would be desirable and this partly explains the continued unavailability of the software components. Preliminary testing has, for instance, indicated potential problems concerning certain graphics cards. Nonetheless, the software has been proven to perform reliably with the MOTU 24 I/O, MOTU 2408, RME Hammerfall, and a few other ASIO-compliant audio interfaces.

Furthermore, although the software has never crashed to date, measures in place to deal with this eventuality should it ever happen are not, perhaps, as extensive as they might be. Although client-side problems are relatively inconsequential, a serious server-side software failure could potentially result in high-amplitude digital noise from all of the audio interface outputs. If this was to happen, and the server machine had stopped responding completely, the only option would be to switch the power off on the audio interface. This is clearly not ideal and further work is therefore needed in this respect. Another minor problem has been identified, and this concerns the cold-booting of the system. Communication with the MOTU 24 I/O is not always successfully established in the first instance, and a re-boot is sometimes needed. However, as the system functions perfectly once it is 'up and running,' it seems most likely that this is an issue relating to the audio interface drivers, and MOTU has been contacted in relation to this.

Software aside, the integration of hardware components – as well as decreasing setup times – improves reliability by minimising the likelihood of equipment being incorrectly patched. With normalised components in place, the system can be rigorously tested once and left in a stable state. In terms of the hardware, there has been one small problem with the control surface, in which one of the 25-pin D-connector sockets (used to connect to the Ircam AtoMIC) was found to be slightly loose. This can almost certainly be attributed to the nuts that secure the connectors to the aluminium casing rattling themselves loose during transit, and could be easily rectified by replacing the nuts with Nylock equivalents.

The modularity of the system – mentioned in most of the criteria given – is also beneficial from the perspective of reliability, because no single component necessarily depends on all of the others to function correctly (obviously this is truer in some cases than in others). This is beneficial in two ways. Firstly, any damaged or faulty components can be replaced individually. Secondly, if a malfunction should occur and it is not possible to repair or replace the component in question, there is a good chance that the system will be able to function in a depleted state (again, this obviously depends on the component it question). Explicit considerations in the design of the control surface, for example, have been made in this respect. Internally, its circuitry is modular, and has been designed so that any defective module can be easily accessed and replaced quickly and easily. If this is not possible, all of the remaining modules will function perfectly. Because each module accommodates four faders, if a module is faulty during a performance, twenty-eight fully functioning faders will still be available. The internal design of the Control Surface is described in Appendix 2.

Overall, it can be concluded that considerable measures have been taken to ensure that the M2 performs reliably. The system may not theoretically be as reliable as some of the other systems described, simply because it is more complex, and incorporates software components. However, there is actually very little material evidence to support this supposition. As with all of the criteria, certain aspects in which improvements could be made have been identified.

5.6.7. Ability to Cater for the Technical Demands of Live Electroacoustic Music

The ability to cater for the technological demands of a wide variety of electroacoustic works is largely attributable to the M2's flexible signal routing and mixing architecture. Audio sources can be connected more or less arbitrarily to the audio interface inputs and routed in software as required during the concert programme. Each piece can have its own system configuration if necessary, and these can be preset and recalled. In this way, all of the audio sources required throughout the programme can be connected to the audio interface at all times, circumventing the need to physically re-patch. The exception to this under the present prototype of the system is if the total number of different audio input source channels in the concert programme exceeds twenty-four. Nonetheless, this is not an insurmountable limitation of the system, as the software natively supports ASIO compliant audio interfaces of an arbitrary number of channels. As described in section 5.6.1, the MOTU PCI card can in fact host up to four 24 I/Os simultaneously, for a total of ninety-six possible analogue input and output channels maximum. This should be ample for all but the very most varied and ambitious concert programmes. Nonetheless it should be recalled that as new technologies and capabilities become readily available, the very nature of the electroacoustic idiom itself – as described in Chapter 2 – is to embrace the new creative possibilities. If ninety-six inputs and outputs are available at present, this may be insufficient in future.

The M2 is also rather more accommodating than some of the other systems described, in terms of its tolerance for works of varying numbers of channels. Again, this is largely attributable to the fact that the signal routing and functionality of the faders can be different for each piece. A stereophonic work can make use of the interface to diffuse pluriphonically to a number of stereophonic coherent loudspeaker sets within the array. Using a different pre-configuration, an octaphonic work could be performed uniphonically with one fader assigned as a level control for each of the eight loudspeakers being used. The next piece again might call for a quadraphonic tape part to be diffused alongside a stereophonically encoded acoustic instrument, and this could also be accommodated, with independent diffusion of the two coherent audio source sets attainable. Technically, the system is able to accommodate each of these distinct demands as consecutive items on a concert programme, with no physical re-patching in between works. If this is to be realised in practice, a carefully planned loudspeaker array will clearly be required, that is able to cater for all of the necessary CLS shapes (see section 3.8.1). This will be discussed further in the next chapter. It will also be noted that the third hypothetical example given (quadraphonic tape with stereophonically amplified acoustic instrument) could not be accommodated by the present system without external microphone pre-amps, and this should certainly be considered as a possible addition to the rack-mount system (even generic systems – the simplest and arguably crudest of them all – normally cater for this).

5.6.8. Ability to Cater for both Top-Down and Bottom-Up Approaches to Sound Diffusion

The ability to pluriphonically diffuse a wide variety of multichannel works, making ergonomic use of the interface in various different ways, is sure to be of considerable benefit to top-down practitioners. The 'wave,' 'chase,' and 'randomise' effects may also prove useful as a means to attaining certain actions that would be difficult to achieve manually.

Although some processes are, to a degree, automated, there is in fact little scope for *accurately controlled* automation within the system; those actions that do occur automatically not being controllable to the extent that certain practitioners might wish. All of the 'effects' are continuously iterating processes and the performer will not know – when fading in a chase sequence, for instance – at what part of the cycle the effect will be when it becomes audible. This has been observed by a number of users of the system and is clearly something that needs to be addressed.

Indeed, few – if any – of the paradigms on offer are 'accurate' in terms of the parameters they manipulate, the system having very much been geared towards perceptually-evaluated *real-time* performance. While this is likely to be of enormous benefit to certain practitioners, what it ultimately amounts to, in respect of the present criterion, is a considerable top-down bias. This is not entirely surprising as the system was developed from within a strongly acousmatic culture at USSS, whose overall preferences tend to swing toward top-down methodologies. If the system is to be widely adopted, however, more attention will need to be paid to the potential requirements of bottom-up practitioners.

It is, nonetheless, true to say that the limitations cited in the present section are merely software constraints and, as such, none are directly attributable to the physical design of the system itself. With even a few relatively minor adjustments to the client software, the M2 has the potential to be equally beneficial to both top-down and bottom-up practitioners alike.

There are additionally a few aspects of the M2 system that do not fall under the given evaluation criteria quite so easily. The ability to run the client and server applications on separate computers is a good example, as this offers a number of fairly unique benefits. It is not always convenient, for example, for the physical signal routing to take place at the diffusion console. This can often result in an unmanageable tangle of cabling around the performance instrument. Because the *SuperDiffuse* server machine performs all of the audio mixing, this can be placed in a location that is convenient for easy and neat cabling to all of the loudspeakers. The exact location is likely to vary from venue to venue, but this does not matter because the two machines can be placed separately. Alternatively, the server machine can be located at a distance from the audience, minimising the impact of noise from cooling fans that has occasionally been raised as an issue. It will be noted that, because communication between the client and server machines is via the TCP/IP, so wireless Ethernet is a possibility. Furthermore, because the client machine does not handle any audio directly, so it can quite feasibly run on a machine of relatively low specification, a laptop for instance. This reduces clutter around the diffusion console, which can very frequently be observed in performances staged using other systems.

5.7. Summary

In comparison with those systems evaluated in Chapter 4, the M2 presents improvements in some important areas. The system is very portable, and therefore comparatively easy to transport between venues. Setup times are also significantly reduced, owing to the integration of key components within the flight case. Setting up the loudspeaker array, however, remains time-consuming, and optimisations (perhaps in the form of rationalised systems for the distribution of mains power and audio signals) could be made in this respect.

In terms of its signal routing and mixing architecture, the M2 is completely configurable, and the settings can be quickly and easily changed from piece to piece by loading preconfigured files. The mix matrix architecture allows any audio input to be dynamically mixed to any audio output, and the internal architecture is governed by the requirements of individual works rather than by the system itself. In this respect, the M2 is more flexible than any of the systems described previously. However, the time taken to configure the system in the first place has been identified as an area in which further optimisations could be made.

The issue of interface ergonomics has been addressed by providing extended functionality via the familiar mixing desk fader paradigm. Again, this can be configured differently for each piece. In this way, improved ergonomics and increased functionality are not at the expense of ease of use, and the system does not deviate radically from familiar conventions. However, as described previously, this has proved at times to be a problematic issue in itself.

Every effort has been made to ensure that the M2 is widely accessible in terms of the technologies required to build it. The modular nature of the system, to a considerable extent, ensures that this is indeed the case: a wide variety of different components (control interfaces, audio interfaces, *et cetera*) can be integrated without affecting the overall functionality of the system. The only real problem in this respect rests with the software, which is currently unavailable outside the University of Sheffield Sound Studios. Availability online would be a fairly straightforward solution to this.

In terms of its ability to accommodate both top-down and bottom-up approaches to sound diffusion, the M2 demonstrates a considerable bias in favour of the former. This is primarily because the nature of control and diffusion techniques offered is more amenable to a perceptually evaluated and 'improvisatory' approach than it is to the precise and consistent articulation of predetermined actions. Accordingly, increased creative scope is offered to the top-down practitioner, with considerably fewer apparent advantages with respect to the bottom-up approach. This would appear to be a difficult problem to address, because it depends so fundamentally upon the holistic nature of the system as a creative framework. Nonetheless, the problem could be addressed in the first instance with relatively minor adjustments to the client software.

In summary, the main advantages of the M2 system are: portability and quick setup times; configurability in terms of mix architecture and diffusion control; configurability in terms of the hardware resources required; increased creative scope and improved ergonomics from a mainly top-down perspective. The main disadvantages are: inadequate support for bottom-up methodologies; time taken to configure the system prior to performance; unavailability of the software. The M2 demonstrates considerable technological and logistical improvements in comparison with many existing systems, with slightly more attention needing to be paid to aesthetic issues (particularly with respect to bottom-up techniques).

Much has been learned through the design, development, and implementation of the M2 Sound Diffusion System, and its subsequent use in staging public performances of electroacoustic music. Many significant improvements in comparison with other diffusion systems have been realised, and the M2 is, in various ways, unique in terms of the creative possibilities it offers to performers of electroacoustic music. There are, of course, issues that remain unresolved, and additional issues have arisen from the unique nature of the system itself. This is only to be expected given what has been observed about the nature creative frameworks in general. Strictly, the M2 cannot therefore be presented as a comprehensive solution to the demands of live sound diffusion, although it can indeed be argued that, in certain important respects, the system comes closer to attaining this ideal than any of the other systems reviewed. If a truly comprehensive solution is ever to be realised (and there is no guarantee that this is even a realistic possibility given the complex and multifaceted demands of live electroacoustic music) then further research is needed. Some specific areas in which further research would be beneficial will be identified in the sixth, and final, chapter.

6. A Way Forward for the Design of New Sound Diffusion Systems

6.1. Introduction

This final chapter will present some overall conclusions and propose some important considerations to be made in the design of future sound diffusion systems. It should be recalled that areas for future research are proposed on the basis of the following (very broad) questions, as outlined in the introduction:

- What can we learn from the technological demands of the electroacoustic idiom?
- What can we learn from the aesthetic nature of electroacoustic music?
- What can we learn from the technical nature of sound diffusion itself, and from differing approaches to and aesthetic attitudes towards it?
- What can be learned from existing sound diffusion systems?
- What can we learn from the design and implementation of the M2 Sound Diffusion System and from the feedback obtained from practitioners who have performed with it?

By adopting an inclusive approach in attempting to answer these questions, and with due consideration given to the evaluation criteria proposed in section 4.3, it is proposed that new systems appropriate for the performance of electroacoustic works from across the technological and aesthetic spectrum of the idiom can be devised. It is hoped that such systems might gain more widespread acceptance throughout (and perhaps even beyond) the electroacoustic community.

What follows is not a design specification, as such, but rather a collection of observations and suggestions that could, conceivably, be developed into one; further refinement of the ideas would certainly be required. This review will begin with a summary of some of the general issues that have been raised, referring back to previous chapters. This will be followed by some more specific suggestions regarding the kind of system that might be required to address these issues and, more specifically still, a few particular ideas and concepts that could potentially be beneficial.

6.2. Some Advantageous Paradigms and Architectures

Through research into existing diffusion systems, and in response to the criteria described in Chapter 4 (as informed by the technological and aesthetic demands of the electroacoustic idiom), the following paradigms and architectures are proposed as useful in the development of future systems.

6.2.1. Matrix Mixing Architecture

Many of the problems associated with existing diffusion systems arise from the inflexibilities often inherent in mixing hardware, particularly that which has not been designed specifically for the purpose of sound diffusion. If the architecture itself limits the ways in which audio input streams can be diffused to audio outputs, this is likely to have a negative impact on both the technological range of works that can be accommodated (evaluation criterion 4.3.1), and on the range of ways in which diffusion can be realised (of particular importance to criterion 4.3.2). Such limitations are in evidence in both of the generic setups, because the architectures of the most common studio mixing desks tend only to easily support two-channel (stereophonic) CASSs. The same can certainly be said of the BEAST, Creatophone, and Cybernéphone systems.

The advantages of the matrix mixing architecture as a solution to this problem have been described previously. The matrix itself can be effectively implemented in hardware or software but, for ergonomic reasons, it is important that software-based *control* of the matrix be made available. Hardware mix matrices (such as the Richmond AudioBox, cited in section 4.14) are advantageous insofar as they do not depend on the processing power of a host computer, and potentially offer increased reliability. However, once built, the specification (number of inputs and outputs, nature of DSP performed, *et cetera*) of a hardware matrix cannot be modified as easily as a software implementation. With computer processing power increasingly less of an issue in recent times, this would appear to be a strong argument in favour of host-based software-emulated matrices, from the perspective of continuously improving developments in functionality at least.

In practical terms, however, there are always constraints. Paradoxically, in its hardware incarnations the matrix mixer is almost certainly among the *least* ergonomic control interfaces for the live diffusion of sound, primarily because of the number of parameters (individual controls) the user has to directly interact with. Unsurprisingly, not one of the systems reviewed features a hardware matrix mixer that must be operated in this way and, as far as the author is aware, no such system is in regular use for the live diffusion of electroacoustic music. It is for this reason that the mix matrix paradigm is best controlled indirectly, via software abstraction.

6.2.2. Software-Based Systems

One of the greatest advantages of software-based diffusion systems is that they easily afford the possibility of highly abstracted control paradigms, as simply exemplified in Figure 47. Using the basic example of four loudspeakers, if the system is based on direct control of output parameters, then one control will be needed for the output level of each loudspeaker. In an abstracted control system, the output level parameters are calculated algorithmically from input (user) parameters in any of an infinite number of possible ways. This is in contrast with strictly hardware-based systems, which are far more likely to necessitate direct one-to-one control over output parameters.

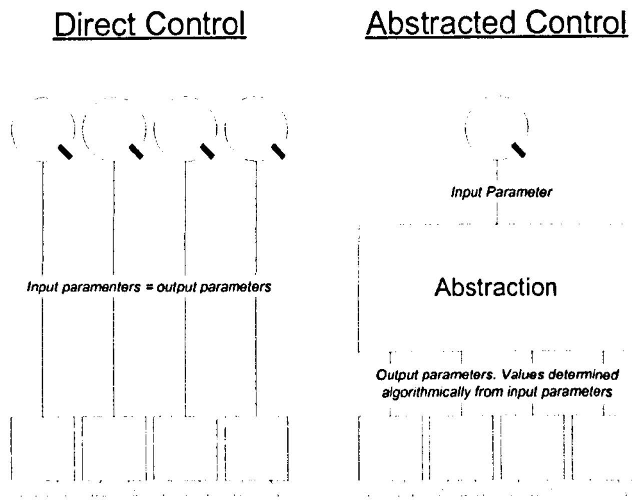

Figure 47. Simple diagrammatic example showing the difference between direct control of parameters versus abstracted control.

Software systems are also advantageous in that developments and improvements can be implemented fairly quickly and economically in comparison with hardware systems, and can be installed on any computer (licensing issues permitting) for increased portability.

6.2.3. Flexible and Accessible Operational Paradigms

Perhaps the most important factor to consider in the present context is the fact that diffusion has traditionally been – and in most cases continues to be – executed with mixing-desk-style faders. Thus, the fader is generally regarded as an approachable paradigm in the context of sound diffusion and, as observed by Harrison in section 4.3.5, many sound projectionists are highly experienced and skilful in its use. This cultural familiarity is extremely beneficial because, on a personal individual basis, it augments the range and quality of diffusion actions that can be realised and, on a more general level, it allows for the raising of performance standards by way of teaching and so on.

As previously discussed, however, the fader – particularly when implemented in conjunction with an overly restrictive mixing and routing architecture – can be fairly non-ergonomic, no matter how familiar it is. Several of the systems evaluated – most notably the Creatophone and BEAST experimental system – indicate an interest in developing and improving this paradigm, however – also noted previously – such explorations are fairly limited in scope. It can also be seen that those systems incorporating non-fader-based control interfaces nonetheless adhere to paradigms that are ostensibly 'familiar' within a musical and/or electroacoustic context. The ACAT Dome and DM-8, for example, incorporate standard keyboard and mouse interfaces that will nowadays be absolutely familiar to every composer of electroacoustic music; the piano-keyboard-style interface of the Kinephone (although in certain respects more specific) is familiar in an even broader musical context.

The use of familiar interfaces in sound diffusion has tangible benefits with respect to this criterion and, it is proposed, such practice should be continued to a certain extent in the development of new systems. However, it is also proposed that more consideration needs to be given to the appropriateness of such interfaces for the specific task of sound diffusion. Familiarity in one context does not guarantee familiarity in another.

New systems that enhance creative scope and improve interface ergonomics, while at the same time *supplementing* existing diffusion paradigms rather than completely replacing them, are required. In this way, a transition towards more flexible interfaces could be made progressively, without alienating practitioners used to well-established paradigms. Some more specific suggestions will be given in section 6.3.4.

6.2.4. Client/Server Architecture

There is much to recommend the development of further client-and-server based systems. This architecture has a number of considerable benefits, not least including the fact that new client applications could be devised without having to re-design the underlying mix engine (whether this be hardware or software). This would be a time-saver for software developers, but also useful in that various client options could be made available to performers, and potentially opened up to third-party development. Specific suggestions, in this respect, will be made in section 6.3.

6.2.5. Scalable Systems

The benefits of a compact and portable system, whose individual components are well integrated, were outlined in section 4.3.7. This raises important questions with regard to the *degree* of portability that can realistically be achieved in a sound diffusion system. Certainly the number of loudspeakers required for even a relatively modest array would seem to conspire against this objective. It therefore goes without saying that a sound diffusion system is very unlikely to achieve the same degree of portability as, say, a laptop computer! Working within the bounds of what is reasonable, however, it is clear that certain optimisations could be made with respect to this criterion that do not seem to have been fully addressed in all of the systems described previously.

A system that is highly scalable is to be recommended, partly because this allows for easy adaptation to a number of different performance scenarios (ranging from full-scale performances with large loudspeaker arrays to smaller performances in intimate venues) and also because it potentially allows for performers to practice their art outside the performance situation using the correct 'instrument.' This latter point would seem, intuitively, to be an absolute prerequisite for any musical performance practice, but – as observed previously – is inadequately catered for in the case of sound diffusion. This is largely owing to potentially impractical technical demands, large loudspeaker arrays being a good example, and some specific suggestions as to how this problem can be addressed (learning from existing diffusion systems) will be proposed in section 6.3.2.

6.2.6. Accessible System Prerequisites

If the skills required to operate non-familiar diffusion paradigms are to be assimilated by practitioners, then this requires a relatively high degree of access to the systems in question. This is doubly true given that the rehearsal time available in concert situations is almost invariably severely limited. On this basis a case can be made for new systems that incorporate standard or otherwise easily available and well-supported components, and perhaps systems that can be built to a variety of specifications without affecting the overall functionality.

6.2.7. Modular Systems

Several of the attributes described above indicate that a modular system is needed. The concept of modularity can apply to both software and hardware systems, and is based around the idea of reducing a system into its component parts in terms of their discrete functionalities.

In abstract terms, a *non*-modular system effectively comprises a single entity or 'black box.' The internal functionalities of that system are therefore closed, and input data (including audio streams and control data, in the case of sound diffusion) are manipulated into output data within a single unitary process. This is illustrated in Figure 48(a), below.

A first-order modular system addresses that unitary process and breaks it down into several constituent processes. In Figure 48(b), we effectively have three 'black boxes' instead of one. The overall functionality is the same but, provided the appropriate communication protocols between modules are observed, modules *A*, *B*, and *C* are can be changed. An insert channel on a hardware mixing desk is a good example: the function of the insert can change (e.g. it could be a reverb or a compressor) but the abstract structure of input-process-output remains the same; the communication protocol is, in this example, transitory-analogue encoded audio streams provided via balanced quarter-inch jack sockets.

A second-order modular system provides a further level of abstraction (an 'abstraction layer') in between each module of the system, as illustrated in Figure 48(c). Basically, an abstraction layer can be regarded as a 'converter' between communication protocols. Say process *A* outputs digital audio data in 16-bit 44.1 kHz format, and process *B* is designed to receive input in 32-bit 48 kHz format. Under scenario (b), the two modules would be incompatible. In scenario (c), however, the abstraction layer between processes *A* and *B* could be used to convert between the two formats and therefore establish correct communication between the modules. This is, basically, how Steinberg's ASIO abstraction layer allows audio interfaces of differing specification to communicate with software.

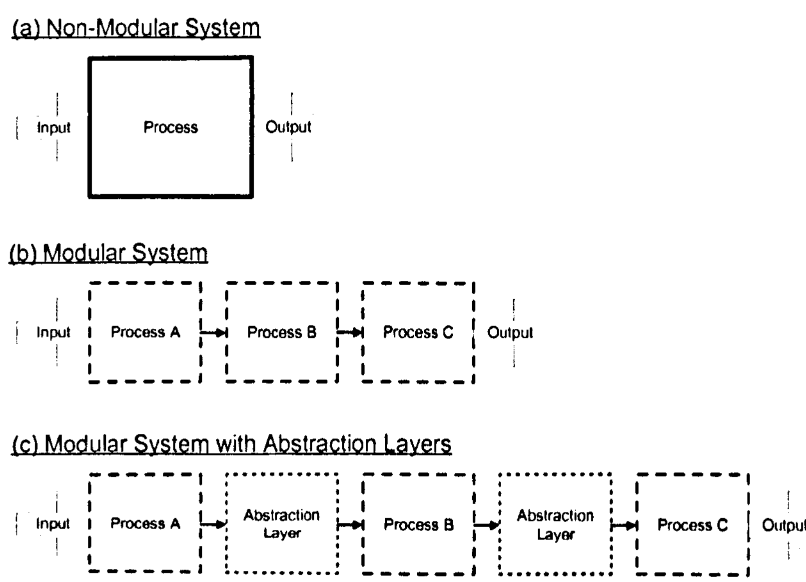

Figure 48. Diagrammatic models of: (a) non-modular system architecture; (b) modular system architecture; (c) modular system architecture incorporating abstraction layers.

Scenarios (b) and – even more-so – (c) are advantageous in that we have access to data at various transitory stages between (overall) 'input' and 'output.' This means that various interventions can be performed before passing data on to the next module. (Comparison with 1.6.1 on page 11 will confirm that this is a very appropriate analogy for the compositional practice of electroacoustic music in general). Modular architectures (the client-server paradigm itself being a good example) therefore have many advantages, most notably that they are potentially hardware independent, that is, the 'same' system can be built with a wide variety of different hardware components. The same applies to modular software designs. Some argue that modular designs can complicate and prolong the design process in certain respects, but it proposed that the advantages with respect to the practice of sound diffusion amply compensate for this.

Overall, it is concluded that due consideration of the criteria summarised in section 4.3 in the design process should result in sound diffusion systems that are more broadly appropriate for the task in hand. Sound diffusion is a unique creative practice with its own particular demands, and flexible systems designed and built specifically for the purpose are required. It should also be noted that a sound diffusion system consists of *all four* of the components identified in Figure 20 (page 168). The flexibility of a system, as observed, can be fundamentally compromised by poor design considerations in any of the four main component areas, and therefore a design approach that applies the evaluation criteria to each individual component, in turn, is to be recommended.

6.3. Some more Specific Possibilities

The following sections will describe some more specific functionalities that will be useful in the development of diffusion systems that are more flexible in terms of the criteria given. Some of these will take the form of potential developments to the M2 system (simply because this is a logical perspective for the author to adopt), while others would require a full re-design in order to be implemented, but these suggestions are broadly intended to be useful for the design of any system.

6.3.1. Increased Functionality via CASS and CLS Concepts

It is suggested that a practical implementation of the CASS and CLS concepts (outlined in Chapter 3) would be useful. This would allow loudspeakers within the array to be specifically grouped by the system in a manner that reflects their prospective use for the diffusion of a particular electroacoustic work. Multiple group membership for individual loudspeakers is possible (see section 3.8.5), meaning that a loudspeaker array can potentially contain many more coherent sets than there are individual loudspeakers. Of course, the CLSs defined within the array are very likely to be determined by the number of channels in the CASS (or CASSs) for that piece (see sections 3.6.5 and 3.8.3). Therefore, a means allowing the user to make this information known (i.e. how audio input channels are to be grouped) to the client application would also be necessary.

The CASS/CLS paradigm has many potentially useful applications, ranging from simple control over multiple signals with one fader (this is effectively already possible in the M2 and some other systems), to more creative and experimental applications. Consider, for example, the possibility of a 'point source panner' designed specifically for panning CASSs to CLSs. In a 'normal' point source panner (such as the one included in Steinberg's Nuendo for the purposes of 5.1 panning[247]) one would expect that the 'position' of the audio source relative to each loudspeaker in the array should determine the relative amplitude of that signal to each loudspeaker. In principle, a CASS-to-CLS panner is the same, only there are restrictions in terms of which audio source (CASS) channels can be sent to which loudspeakers. Specifically, a CASS channel can only be sent to a loudspeaker if that particular input-to-output routing is appropriate (whether or not this is the case could be, to a greater or lesser extent, user defined). As an example, consider the scenario illustrated in Figure 49 (below). Here, we have a 'main eight'

[247] Steinberg (2005). "Steinberg". Website available at: http://www.steinberg.de/Steinberg/defaultb0c4.html.

array; we will assume that we are dealing with a two channel source (i.e. one stereophonic CASS consisting of 'left' and 'right') and have therefore decided to treat the loudspeaker array as four stereophonic CLSs. If we have a user interface control that allows us to define the 'position' of the CASS along an imaginary axis running from the front to the back of the venue, then that CASS can be dynamically 'cross faded' between the four CLSs with a single control. This paradigm could be described as a 'multi-point cross-fader': with the fader at 'maximum' position, the CASS emanates from the 'distant' CLS; with the fader at 'minimum' position, the CASS emanates from the 'rear' CLS; at intermediate points, interpolation between CLSs occurs. The control would not necessarily have to operate on a simple linear basis, but could have custom interpolation functions applied if necessary.

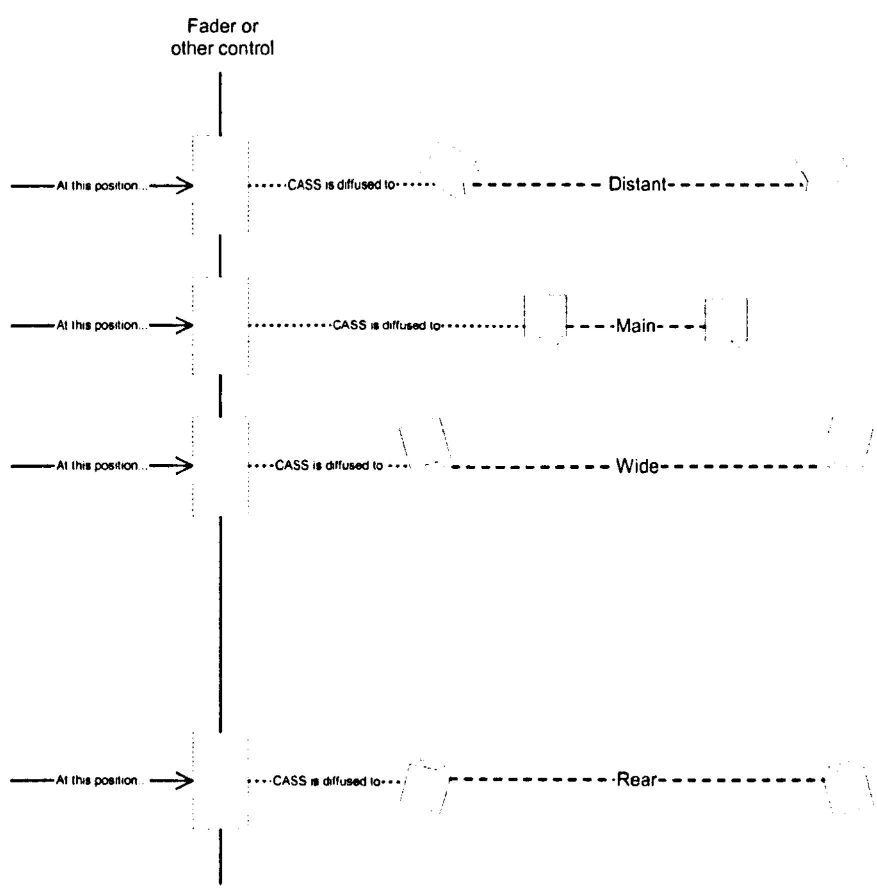

Figure 49. Diagram illustrating the operation of a 'multipoint crossfader' as a means to dynamically diffusing a CASS to multiple CLSs.

Of course, the scenario described above would not be sufficient on its own, but rather could form part of a repertoire of useful control paradigms, or diffusion 'actions.' Overall, a practical implementation of CASS and CLS concepts would allow performers to deal with their audio source channels in a more coordinated and efficient way (rather than individually) and would also represent a step towards establishing interface/loudspeaker-array independence in diffusion systems.

6.3.2. Interface/Loudspeaker-Array Independence

Any diffusion system that operationally depends on our knowing the exact nature of the loudspeaker configuration to be used will present difficulties in the use of smaller arrays for practice purposes. This is a tricky problem to address, but much can be learned, in abstract terms, from the Ambisonic system (described in section 4.13), which – given identical input data – is capable of producing (ostensibly) the same spatio-auditory results over a variety of different loudspeaker arrays. Although Ambisonics is of questionable widespread suitability for the purposes of sound diffusion, conceptually, this particular attribute would appear to be very useful. The DM-8 system (section 4.14) also offers this functionality.

One problem with both of these systems, however, is that they impose specific restrictions with regard to the shape of the loudspeaker array. Ambisonics, for example, requires loudspeakers to be arranged in opposite pairs on the surface of an imaginary sphere, and focused on the central point. Similarly, the DM-8 will allow for the automated performance of pre-planned diffusions to be realised via a variety of loudspeaker arrays, provided that they are arranged in a circular formation and, again, equidistant from the central point. It is true to say that, in both cases, the circular/spherical formation is primarily a function of mathematical convenience, and is also due to the fact that both systems are designed with a degree of 'precision' in mind (in this respect, they are both bottom-up systems). There is, of course, no reason why such capabilities should not be incorporated into future systems, but it is argued that any system offering *only* this kind of functionality would only be of use to a fraction of the electroacoustic community.

Such an approach could be supplemented by a parallel system in which loudspeakers are arranged graphically in the client application, reflecting their actual deployment within the venue. Loudspeakers would then be grouped into coherent sets (as suggested in the previous section) according to the work in question, and then further defined in terms of 'higher level' criteria. A system of common categories, some of them arranged hierarchically, would be devised. Further research would be needed for this, but as a suggestion, categories relating to the 'higher level' positioning of loudspeakers (attributes other than position might also be considered) could include 'mains,' 'wides,' and so on, to use BEAST terminology. A sub-category of 'wide' might be 'extra-wide,' for instance. User-defined categories might also be a useful possibility. In this way, the diffusion client would 'know' about the positions, groupings, and prospective uses of the different loudspeakers in the array,[248] and this information could be used to 'map' between different loudspeaker arrays according to the nearest appropriate match. Clearly, this would only be possible using a software-based system. A hypothetical example is given in Figure 50. Loudspeaker arrays (a) and (b) are clearly related – either (a) is regarded as a 'scaled-down' version of (b), or (b) is regarded as an 'extended version' of (a) – but this conclusion is most easily reached on the basis of *perceptual* evaluation. Loudspeakers are all arranged into stereophonic CLSs. CLSs that can be regarded as categorically 'similar' between the two arrays are colour coded accordingly. Of course, this could be done in a number of different ways according to the requirements of the piece and the

[248] In most systems – including the M2 – the diffusion client or equivalent does not 'know' anything (or 'care,' if this is a better personification) about the loudspeakers, and refers to them only in terms of the audio output to which they are connected. This direct approach is a 'bottle-neck' in the present context, resulting in the array dependence that we are presently trying to evade. Note that the Ambisonic and DM-8 systems circumvent this by applying a further level of abstraction. The presently suggested approach is essentially similar but less geared towards 'mathematical precision' and more towards 'perceived results.' In this respect it is presented as potentially useful to top-down practitioners in particular, and should probably be implemented along-side more bottom-up techniques.

preferences of the performer. The groupings might be substantially different for a 5.1 CASS for instance. If such judgements are to be useful in a software context, mechanisms must be in place for the software (which, under present technological constraints, only understands objective parameters) to have access to this information, and clearly this must be provided by the user. In short, it would be useful to be able to refer to loudspeakers – and groups of loudspeakers – without referring directly to the audio interface outputs to which they are connected. It is proposed that such an approach could potentially be developed into a solution to the problem of scalable loudspeaker arrays that is more useful to top-down practitioners in particular.

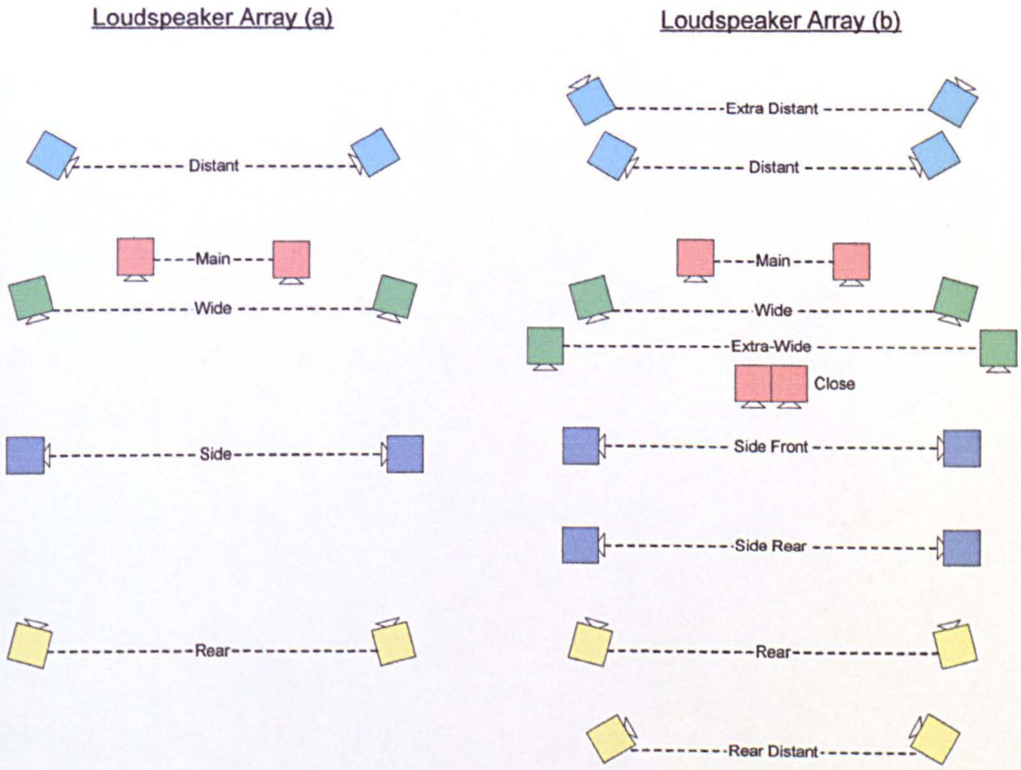

Figure 50. Two loudspeaker arrays that can be recognised (perceptually) as related.

6.3.3. Diffusion 'Actions'

Implementation of items 6.3.1 and 6.3.2 in particular would be a step towards establishing platforms within which performers could diffuse works more directly in terms of the specific diffusion 'actions' involved. Present systems mainly afford the user direct control over loudspeaker amplitudes, or else slightly more abstract control over the 'position' of sound sources (recalling the ACAT Dome and DM-8 examples). However, these are not always obviously amenable to the kinds of actions that practitioners wish to perform, which are often more easily associated with 'higher level' descriptors. The 'effects' provided by the *SuperDiffuse* client application (see section 5.4.2.3) represent initial steps towards remedying this, and this could be developed more thoroughly. A set of highly configurable and intuitively controllable (see section 6.3.4) diffusion actions based loosely around Smalley's taxonomies ('divergence,' 'convergence,' 'contraction,' 'parabola,' *et cetera*)[249] would surely be of interest to top-down practitioners in particular. For the bottom-up practitioner, more objectively satisfactory actions would be desirable, perhaps with more of a focus on quantifiable processes, and with accurately controllable parameters. Both approaches, however, could obviously be incorporated into the same overall software framework. Of course, diffusion actions could be taken to include extended DSP functionality in addition to simple amplitude manipulation.

6.3.4. Support for Multiple Physical Control Paradigms

Control over the output level of a single loudspeaker can easily and ergonomically be achieved with a single fader. Simultaneous control over the levels of a group of loudspeakers can also be ergonomically achieved with a single fader, so long as real-time control over the relative amplitudes is not required. This method is effective because the control paradigm is appropriate to the action being performed. When diffusion actions become more abstract than this (and the previous three sections have largely been devoted towards suggesting that this should at least be an option), it is certainly time to begin questioning the

[249] Smalley (1986). "Spectro-Morphology and Structuring Processes". In Emmerson [Ed.] *The Language of Electroacoustic Music*: 74.

appropriateness of the fader as a control paradigm. For this reason, an exploration of alternative paradigms *in addition* to familiar ones is suggested. These might include joysticks, foot-pedals, buttons, switches, track-balls, and any number of other input devices. The idea is that the performer, or performers, can choose how to map these controls onto the diffusion actions they wish to perform. Ideally, an architecture that abstracts the specific control method from the means of communication of control data to the diffusion client is required. Technically the Ircam AtoMIC is already used by the M2 system for this purpose, but this is neither the most elegant, nor economic, nor user-friendly solution. What might be better is a control surface designed specifically with a modular approach to physical control paradigms in mind. A device in which 'control modules' are interchangeable might be desirable, and a sketch is given in Figure 51, below. Internally, the current M2 control surface is modular in architecture (see Appendix 2) and could therefore be adapted to suit relatively easily.

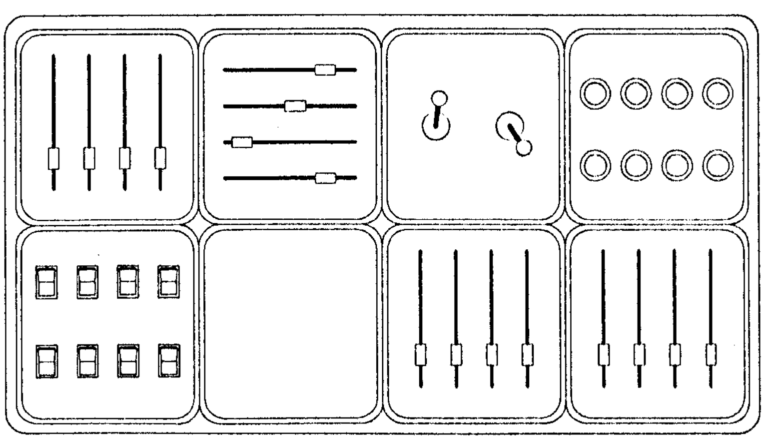

Figure 51. Sketch of a fully modular Control Surface design with interchangeable control units. Modules shown are (left-to-right, top row first): vertical faders, horizontal faders, joysticks, momentary switches, switches, dummy panel for use when no module is installed.

6.3.5. Extended DSP Functionality

Many diffusion systems – including the M2 – only really allow us to perform amplitude manipulation on the signals. In some cases this may be sufficient, but there are situations in which extended functionality may be necessary or desirable. The Cybernéphone (section 4.11), for example, affords performers access to time-delay, filtering (EQ), phase adjustment and pitch shifting functionalities. With the possible exception of the latter, all of these would be of definite use to bottom-up practitioners as means to 'correcting' less-than-ideal loudspeaker arrays. Applications for top-down performers would probably adopt a more creative guise. This gives an indication of some of the factors that must be considered in the design and implementation of high-level diffusion actions incorporating extended DSP functionality.

6.3.6. External and/or Automated Control via Third-Party Software Development

An advantage of modular software-based diffusion technologies rests in the fact that the appropriate interfacing methods and communication protocols can be released to third-party developers (or directly to practitioners) in the form of software development kits (SDKs). The benefits of this approach are well-proven in the commercial sector. Steinberg's VST (Virtual Studio Technology) and ASIO (Audio Streaming Input and Output) software development kits – both freely available to third-party developers on request – have resulted in the development of a wide variety of compatible applications (some of them taking the shape of plugins), many of which are freely downloadable.

A software development kit for the purposes of sound diffusion could take any of several possible forms. Using the client-server model as an example, it would be possible to allow third-party software to communicate directly with the server (in the case of the M2, this would allow direct control over the mix matrix, effectively allowing developers to write alternative client software), or with the client (third parties would essentially be developing a 'front end for the front end,' and would be able to automate the client software for instance). Perhaps both approaches could be adopted.

One approach that can already be realised with the current M2 system is the provision of MIDI control data from a source other than the control surface. The use of Max/MSP for this purpose would open up an entirely different range of possibilities for algorithmic control of the diffusion (as opposed to real-time control via a physical interface), and would provide a framework more suitable perhaps for bottom-up practitioners. Another approach still – suggested in a personal email to the author from Charlie Richmond, Managing Director of Richmond Audio Design – would be the use of the *SuperDiffuse* client application in conjunction with the Richmond AudioBox hardware matrix mixer.

The possibilities are numerous, because any system that exhibits second-order modularity (see Figure 48(c)), will – by its very design – provide several potential 'points of entry' for external intervention, any or all of which could be exploited by making the relevant methods and protocols known to third parties.

6.3.7. Diffusion System Software Plugins

A system for high-level diffusion 'actions' (whether based around amplitude-only diffusion, extended DSP, or a combination) could be plugin implemented. Much like VST plugins, this would allow third-party developers to define their own routines within a standardised framework. This would clearly be advantageous because practitioners would be able to devise routines that are appropriate for their own particular purposes, and in doing so, extend the overall repertoire of diffusion actions on a more general level.

6.3.8. Logistical Functionalities

As described in previous chapters, the task of staging public performances of electroacoustic music can be a substantial undertaking. Research has indicated a few mechanisms that could be put in place to make the task quicker, easier, or otherwise more reliable.

In section 5.6.6 it was noted that contingencies to deal with the event of software failure are not as thoroughly in-place in the M2 system as would be ideal. Where control of signals sent to loudspeakers is governed entirely by software, this is always likely to be a problem unless some kind of in-line attenuation system is provided in hardware. This would purely be an emergency measure, with no diffusion function as such, and could therefore be highly ergonomic. No individual attenuation of loudspeaker levels would be required, and so a design incorporating a single attenuator would be ideal.

An in-built test signal facility would also be beneficial. In conjunction with abstracted signal routings this would allow loudspeakers to be connected to audio interface outputs – literally – arbitrarily, and the relevant mappings made in software, via test signal calibration.

Also in relation to loudspeakers, it has previously been noted that setting these up frequently represents the majority of the setup time overhead in the staging of electroacoustic music concerts. This is often because venues provide relatively few mains outlets, which, in addition, are not always conveniently located. An efficient power distribution system to deal with this would certainly be extremely useful. Because active loudspeakers invariably need to be provided with mains power and signals, there is perhaps also a case for attempting to integrate these two distribution systems in some way, although care would obviously have to be taken to avoid unwanted electrical interference.

6.3.9. Increased Accessibility via Online Resources

If a system partially based around software components is to be implemented, it is very important that these are made readily available, either commercially or freely, and are be well supported in terms of documentation. The internet would seem to be an extremely useful mediator in this respect. Furthermore, if TCP/IP is to be implemented as a means of communication between components in future systems (as it is in the M2), this also means that network facilities would be natively supported by the system, offering the potential for direct software and/or firmware updates and so on. Networking technology – for both creative and purely functional purposes – is becoming increasingly common in electroacoustic music,[250] and this would appear to be an argument in favour of its inclusion in future diffusion systems.

6.4. Areas for Further Research: an Inclusive Approach...

General architectures and specific suggestions aside, this research has additionally indicated various areas in which further research would be beneficial. It seems apparent that the global development of sound diffusion systems is still very much engaged in a research process that is somewhat specific to the particular concerns of the researchers in question. If the approach advocated by this thesis is to be widely implemented, then further research is needed. Research into the ways in which specific top-down and bottom-up techniques can be accommodated into the design of future diffusion systems – in addition to an exploration of how these needs can be accommodated by control interfaces – is particularly important. Overall, this thesis proposes an inclusive approach to the research process, which can be described in terms of the following aspects:

[250] Nicholas Melia and Matt Rogalsky's *Treatise on the Economy and Containment of Despair* (which was performed at the Sonic Arts Network's *Expo966* conference in 2005) is described in programme notes as 'a piece for guitar, violin, and networked electronics.' Sounds encoded via microphones in real-time from the acoustic instruments are streamed across a network to a remote location (in this case the network connection was between Scarborough, UK and Kingston, Ontario, Canada) where they are processed and streamed back to the performance space for diffusion.

6.4.1. ...in terms of Technology

In Chapter 1, an approach was made towards gaining a basic holistic understanding of the scope of electroacoustic music in technological terms. This was summarised in Table 2 (page 23), which comprises six 'base unit' technological profiles. A work of electroacoustic music might make use of one or more of these profiles, in any combination, possibly incorporating all of them. This will, obviously, have an impact on the kinds of demands that will be exerted on a diffusion system if it is to be used for the performance of the work in question. In adopting an inclusive approach it is proposed that a diffusion system should be able to cater for all of these eventualities in the context of live performance.

It could, of course, be argued that this is already the case, and to a certain extent this is indeed true. Most diffusion systems are able, at some level, to accommodate all of the technologies discussed. However, it can be counter-argued that this is merely a fortunate by-product of the fact that widely accepted standards exist for the communication and storage of encoded audio streams. Studio mixing desks tend to have analogue audio inputs: it is therefore possible to connect any device capable of outputting transitorily encoded audio streams. Problems very quickly become apparent, however, when works that vary in technological profile are juxtaposed in concerts. When this is the case, this indicates that the diffusion system is not as well-suited to the task in hand as might be desirable. This is because most applications of mixing desks do not exert these kinds of demands, and it is therefore not usually necessary to cater for them in the design of technologies. A more focussed approach, geared towards the specific requirements of electroacoustic music and sound diffusion – bearing fully in mind that electroacoustic music tends to involve the appropriation of technologies in fairly unique and detailed ways – is therefore to be recommended.

6.4.2. ...in terms of Aesthetics and the Aesthetic Affordances of Creative Frameworks

Top-down and bottom-up attitudes were introduced, and their salient characteristics abstractly defined in Chapter 2. In this chapter it was also noted that creative frameworks – and in the context of electroacoustic music this includes both the means of composition and the means of diffusion – can be 'directionally biased,' that is, owing to their very nature, they can individually exert certain pressures towards either top-down or bottom-up approaches to their use. Some frameworks are, of course, broadly conducive to both approaches, encoded audio streams – and, by extension, the electroacoustic idiom itself – being excellent examples.

A diffusion system should, equally and without prejudice, be able to accommodate both top-down and bottom-up approaches to sound diffusion. This is a tricky issue to approach because it is inherently paradoxical. The most effective means to composing works of a bottom-up nature are likely to be those means that are intrinsically suited to that particular way of working. Conversely, those frameworks that exert a bias in favour of top-down approaches to their use are likely to be more appropriate for the composition of top-down works. The very relevant example of real sounds 'versus' synthesised sounds, as those frameworks that respectively engendered the praxes of *musique concrète* (top-down) and *elektronische Musik* (bottom-up), was given in section 2.3. In the context of *performing* works, the same can be said of sound diffusion systems: top-down works will be most effectively diffused via a top-down framework, and bottom-up works are likely to be more faithfully communicated via a bottom-up framework. Viewed from the perspective of diffusion system design, what this seems to imply is that any design features that are favourable to bottom-up practitioners (those that will result in the better communication of their works) will be correspondingly *un*favourable to top-down practitioners, and *vice versa*. However – and herein lies the paradox – it is very likely that a diffusion system will have to accommodate both kinds of work. Accordingly, the most commonly employed diffusion systems – the generic templates described in section 4.6 – are relatively neutral in terms of their directional bias, but as a consequence are far from ideal from both top-down and bottom-up standpoints.

Two ways in which this frustrating and paradoxical problem could be approached seem evident. Either 'top-down sound diffusion systems' and 'bottom-up sound diffusion systems' are to be regarded as two separate and mutually exclusive frameworks, each with their own unique characteristics and design considerations, or else attempts to develop systems that are thoughtfully designed, highly flexible – and therefore able to cater for both of these contrasting attitudes – are to be made. This thesis opts for the latter choice because it engenders both diversity and an inclusive attitude towards the electroacoustic idiom. The former approach, however, is frequently evidenced and has been observed, and criticised, by Landy:

> One part of the hypothesis introduced above concerns the amount of individual/small group effort going into development and scholarship (as well as composition, in fact) in isolation. [...] Individuals [stake] their claim to an idea, an approach or some such often without adequate contextualisation, but more importantly here without adequate or any feedback or consistent correlation, using methodologies that are often self-referential. Even if we accept that after fifty years we continue to be participants in an experimental phase or process, the implications of continued pockets of isolation are dissatisfying. Our [...] enclosed, self-supporting structure is as endangered as the opera companies and orchestras of the world waiting for their arts [...] subsidies to diminish if not pass away. In other words, if the work of many musicians and [...] researchers remains marginalised, how can people expect all forms of support to continue? The island mentality represents an archaic aspect of academe which hopefully will evolve a little in the coming decades...[251]

Landy's observations raise further important issues with regard to the continued relevance of the electroacoustic idiom in an altogether broader cultural context; we will return to this shortly. In terms of sound diffusion (which is, of course, central to the appreciation of electroacoustic music in a wider context anyway), systems providing

[251] Landy (1999). "Reviewing the Musicology of Electroacoustic Music - A Plea for Greater Triangulation". *Organised Sound*, 4(1): 63, 68.

frameworks that are conducive to both top-down and bottom-up approaches are needed, and in order to achieve this, further research into the various directional biases of prospective techniques and paradigms will be useful.

At present, it can be argued that practitioners can be so strongly biased in favour of their own particular methods, techniques, systems, and so on, that they fail (or even flatly refuse) to acknowledge the validity of other approaches. This is clearly counter-productive in a broader context. It is not by any means suggested that all practitioners should adopt a single mutual and unanimous perspective, but rather that a wider variety of approaches should be regarded as equally valid despite their inevitable differences.

6.4.3. ...in terms of Specific Methodologies and Techniques

In Chapter 3 it was observed that top-down and bottom-up attitudes tend to foster markedly different approaches to the task of sound diffusion. Again, this poses difficult challenges in the design of diffusion systems and – once again – an inclusive approach to this problem is proposed. More research into the specific techniques favoured by top-down and bottom-up practitioners, and with respect to how these can be accommodated into the design of future diffusion systems is needed, and it would appear that a more holistic and inclusive ethos needs to be adopted in order for this to happen. This proposal might seem very similar to that made in the previous section. The distinction is this: in the previous section, research with a view to ascertaining which creative *frameworks* are top-down, and which are bottom-up, was advocated; in the present section, a more in-depth assessment of the different *techniques* (actions) carried out by top-down and bottom-up practitioners during the process of sound diffusion is recommended. These two distinct areas of research would appear to be reciprocally beneficial. If we know, broadly, what top-down and bottom-up practitioners 'want to do' (Table 4 on page 162 – it is hoped – will be a good starting point) and which frameworks will be most appropriate for the facilitation of these actions, then we will be in a position to implement designs that cater both equally and more flexibly and efficiently for both parties.

6.4.4. ...in terms of Potential Applications Outside the Electroacoustic Idiom

Of course, there is no need for advanced sound diffusion systems to be devoted exclusively to electroacoustic music. There are other situations in which the live broadcast of sound via multiple loudspeakers may be necessary or useful. Theatre applications are a definite possibility, and these represent a distinct target market for the Richmond AudioBox.[252] The director of the University of Sheffield Drama Studio has expressed an interest in the M2 system, stating that it would provide tangible benefits in the field of theatre sound.[253] Installation in night-clubs and other music venues might also be a possibility. If such applications are to be considered, this will obviously need to be taken into account at the design stage. Applications with commercial prospects (electroacoustic music is, after all, a fairly specialised area) may be an effective means to bringing much-needed development funds into the design of new systems – as well as offering the potential to raise the profile of electroacoustic music – and should therefore be embraced.

6.5. Summary

What these suggestions ultimately amount to is the endorsement of an altogether more holistic approach to the task of sound diffusion that takes into account all of the technologies involved (both at the present time and in the foreseeable future), different attitudes towards electroacoustic music and the appropriation of creative (technological) frameworks (which, for simplicity, can broadly be characterised as top-down and bottom-up), and differing approaches to sound diffusion in terms of the specific actions taken

[252] Harmonic Functions (2001). "The AudioBox Disk Playback Matrix Mixer". Website available at: http://www.hfi.com/dm16.htm.
[253] David Moore, co-developer of the M2 system. Personal communication with the author.

(again, top-down and bottom-up). There have additionally been suggestions made with a view to making the logistical task of staging electroacoustic music concerts more efficient.

In technological terms, a more abstract approach to the diffusion process would be beneficial. Rather than conceptualising diffusion as a practice in which 'audio inputs' are mixed to 'audio outputs' by moving faders, it would be better to imagine the process as one in which 'sources' are diffused to groups of loudspeakers via a repertoire of 'actions.' The CASS and CLS will be useful concepts if this approach is to be adopted. Alternative control interfaces should also be considered, bearing in mind that each of these is a creative framework in itself, and can therefore have a directional bias.

In aesthetic terms, we must consider the fact that top-down and bottom-up attitudes exist within the electroacoustic idiom. Attempts to reconcile these two opposing philosophies have thus far proved unsuccessful and therefore – for the time being, at least – their respective methodologies must co-exist. This is problematic because the two ideals are, in many respects, antithetical. Nonetheless, through research into the directional biases of creative frameworks, it will be possible to ascertain which means are appropriate for which ends, and how these can be integrated within a system such that top-down and bottom-up techniques are equally catered for.

In short, the design of improved sound diffusion systems assumes an awareness of the kinds of actions performers wish to carry out, and an understanding of the underlying musical beliefs that engender these methodologies. We need to know what kind of musical communication is sought, and what functionalities will be required to facilitate this. Once this has been ascertained, attempts can be made to discover the best means to carrying out those actions in an appropriate way. This thesis has attempted to address these issues, and much ground has been covered, but no single study can be comprehensive. A collaborative approach is therefore needed if new sound diffusion systems are to be beneficial to both individual practitioners and to the electroacoustic community as a whole.

Conclusion

Electroacoustic music is a diverse and multifaceted idiom that is characterised by the appropriation of certain creative frameworks (see section 1.9) in unique and specific ways. The creative frameworks themselves were summarised in section 1.7; these include (but are not necessarily limited to): audio encoding technologies; recording and playback technologies; processing technologies; synthesis technologies; and audio decoding technologies. Software technologies are often used to emulate one or more of these, and can also be appropriated as creative frameworks in their own right. All of these frameworks, in some capacity, deal with encoded audio streams (section 1.6), which can therefore be regarded as a partial defining characteristic of the electroacoustic idiom. The loudspeaker (or, more abstractly speaking, the audio decoding technology) can also be regarded as a defining characteristic because all of the other frameworks rely on it as a means to realising the final auditory result. As far as the performance of electroacoustic music is concerned, there are certain 'instrumentations' (combinations of creative frameworks) that are culturally identifiable as characteristically 'electroacoustic.' These were summarised in Table 2 (page 23), and sound diffusion systems should ideally accommodate all the technological permutations.

Typically, the creative process in electroacoustic music focuses directly, and as the *primary object* of artistic exploration, on one or more of the affordances uniquely offered by the creative frameworks employed. This is as opposed to using the frameworks as a functional means to realising and/or communicating objectives whose primary creative drive has its origins in the affordances of other frameworks. This differentiates electroacoustic music from artistic praxes (such as 'popular' music, for example) that engage the same frameworks in a secondary manner, and also from purely functional applications of the same technologies, although in the former case there is clearly the potential for ambiguity (certain popular music traditions such as 'intelligent dance music' for example). While the creative frameworks in general cannot be regarded as defining characteristics of the electroacoustic idiom of themselves (they have too many other applications), the combination of the frameworks and the unique ways in which they are appropriated *can* be regarded as a defining characteristic.

Within the scope of this definition, approaches to the composition and performance of electroacoustic music can be regarded as either top-down or bottom-up. These terminologies do not connote specific techniques but, rather, broad philosophies with respect to the nature and 'purpose' of electroacoustic music. Of course, specific techniques are invariably born out of these underlying convictions, but the exact nature of these can vary considerably. It is therefore helpful to focus on the general as opposed to the specific: a summary of the opposing traits of top-down and bottom-up philosophies is given in Table 4 (page 162). Historically, this binarism can be observed in the praxes of *musique concrète* and *elektronische Musik*, respectively, and is still strongly in evidence. These contrasting ideologies fundamentally impact upon the nature of compositions (as described in Chapter 2) and the approaches taken in their performance (as described in Chapter 3).

Sound diffusion is the *active process* by which works of electroacoustic music are performed to an audience. As such, this process necessarily involves the use of loudspeakers to decode the encoded audio streams and render the real auditory results. Because multiple encoded streams are very often (in fact, almost invariably) used collectively to encode spatio-auditory attributes (see section 1.6.2), this process can involve the spatialisation of encoded material that already contains spatial information. It seems logical that attempts should be made to preserve this spatial information during diffusion, and it is primarily for this reason that the concepts of the coherent audio source set (CASS) and coherent loudspeaker set (CLS) have been proposed; these were explained in sections 3.5 and 3.7, respectively.

The essential nature and ultimate objectives of the process of sound diffusion vary depending on aesthetic directionality, that is, 'top-down diffusion' and 'bottom-up diffusion' can be regarded as two fairly distinct performance practice conventions. Stemming from the fact that top-down practice is, essentially, a perceptually motivated exercise, the ultimate purpose of 'top-down diffusion' is to present the electroacoustic work in a manner that is perceptually effective in the given performance context (this can include the venue, audience, system, the piece itself, and so on). The efficacy of the results is evaluated via a continuous process of subjective judgement at the time of performance. For the bottom-up diffuser – who is more objectively motivated – the aim is to faithfully communicate the abstract concepts and structures embodied within the given electroacoustic work. In this case, the efficacy of the results is ensured via careful preparation of the performance context (which, again, can include the venue, audience, system, the piece itself, and so on) *prior* to performance. In terms of technique, top-down diffusers tend to favour pluriphonic *interpretations* of the electroacoustic work, whereas bottom-up diffusers are more inclined to prefer uniphonic *presentations* in controlled acoustic environments.

For the top-down practitioner, the encoded audio stream (a defining characteristic of the idiom) is an abstract representation of the perceptual experience of a real auditory event. For the bottom-up practitioner, the stream is a carrier of objective data that collectively express abstract micro- and/or macro-structures. In both cases, it is a fundamental aim of sound diffusion to present the CASS (which, of course, simply consists of one or more encoded audio streams) in a manner that makes its 'contents' apparent to the audience: in this respect, top-down and bottom-up practitioners differ only in terms of what they consider those 'contents' to be (perceptual or conceptual). It can therefore be said that maintaining the coherency of the CASS (or CASSs) is a unanimous objective of the process of sound diffusion.

Sound diffusion systems are the means by which the process of sound diffusion is carried out, and are therefore also the means by which electroacoustic music is performed. As such, diffusion systems should ideally accommodate the full range of technological and aesthetic diversity embodied within the electroacoustic idiom. This includes the purely technological requirements, conditions determined by top-down and bottom-up philosophies, consideration of the practical and logistical nature of the process of sound diffusion itself and, of course, the unique requirements of individual works and practitioners if these deviate from the trends discussed. This is a complex and challenging set of demands indeed, and in Chapter 4, a system of criteria was described, taking all of these factors into account. These are as follows:

- Ability to Cater for the Technical Demands of Live Electroacoustic Music
- Ability to Cater for both Top-Down and Bottom-Up Approaches to Sound Diffusion
- Flexibility of Signal Routing and Mixing Architecture
- Interface Ergonomics
- Accessibility with respect to Operational Paradigms
- Accessibility with respect to Technological System Prerequisites and Support thereof
- Compactness / Portability and Integration of Components
- Reliability

All of the criteria are strongly interrelated, and were explained fully in section 4.3.

On this basis of these criteria it can be stated that current sound diffusion system technology – as described in Chapters 4 and 5 – does not fully cater for the requirements of the electroacoustic idiom. In many cases this is because technological restrictions are imposed, but it is equally often because the systems are inherently biased in favour of either top-down or bottom-up performance practice methodologies. The design of improved systems is therefore necessary.

In designing new systems, reference to the set of criteria outlined in Chapter 4 will be helpful. Consideration of the relative advantages and disadvantages of existing systems is also to be recommended. Modular systems with a high level of abstraction between components are advantageous because they offer technological flexibility and can be appropriated in various different ways. Software-based systems are also worth considering as these offer a high degree of dynamic configuration. At present, the concepts of the CASS and the CLS are maintained mainly by convention: when a diffuser raises two mixing desk faders at once, for example, he or she is effectively taking measures to ensure that the coherency of the stereo CASS is maintained. Logistically, the practice of sound diffusion could be made much more efficient if new systems allowed for the explicit specification of CASSs and CLSs by performers. In this way, the process of sound diffusion could deal with logical *groupings* of encoded audio streams – as is conventionally the case anyway – as opposed to manipulating encoded audio streams individually. This paradigm would also allow for the development of a performance practice that is more oriented towards high level diffusion 'actions' (see section 6.3.3) and less focussed on the individual mixing of audio inputs to audio outputs. This will be particularly advantageous in implementations of the mix-matrix architecture, which – as discussed in section 6.2.1 – is the most appropriate paradigm for sound diffusion but, paradoxically, also the most difficult to control.

Numerous further suggestions were given in Chapter 6; many, many different approaches are possible – and it has not been the purpose of this thesis to state which specific approach is 'best' – but one must always be aware of the fact that sound diffusion systems are creative frameworks in themselves. Because every creative framework can exert a directional bias in favour of top-down or bottom-up approaches to its use, particular care must be taken to ensure that new systems are acceptable to both parties. Clearly, this will have an impact on all aspects of the design, but particular attention should be paid to the operational capabilities of the system and the nature of the control interface. Above all, the design of diffusion systems should always afford careful consideration to the very broad technological and aesthetic nature of the electroacoustic idiom itself.

Appendix 1: An Interview with Professor Jonty Harrison

This appendix presents a verbatim transcription of an interview conducted by the author with Professor Jonty Harrison. The interview took place at the University of Birmingham on 8th September 2004, and has been reasonably often cited in the present thesis. At the time of writing, this interview has not been published elsewhere.

Professor Harrison is among the best established practitioners of live sound diffusion in Europe, if not world-wide, and is the founder of the Birmingham Electro-Acoustic Sound Theatre (BEAST), whose current diffusion system and recent experiments in the field of sound diffusion have been documented in sections 4.8 and 4.9 respectively. The former can effectively be regarded as having set the standard for sound diffusion systems (certainly in the United Kingdom) and is – at the time of writing – currently involved in a European tour. Harrison is, equally, an internationally acclaimed composer of acousmatic works, including the frequently performed and much-cited *Klang* (realised in 1982).

The present thesis owes a great deal to Harrison's published writings on electroacoustic music and sound diffusion, his writings on 'organic' and 'architectonic' structuring principles (see section 2.5.3) having been particularly influential. With reference to Table 4 (page 162), it should be reasonably obvious that Harrison's ethos falls more obviously into the top-down (or, to use his own terminology, 'organic') category than the bottom-up.

Obviously, an enormous debt of gratitude is owed to Professor Harrison for providing this interview, which has been much appreciated by the author. It is hoped that the inclusion of this transcription will both clarify and supplement some of the salient aspects discussed in the main body of the thesis.

James Mooney: As one of the best established practitioners of the art, what does 'sound diffusion' mean to you?

Jonty Harrison: Thank you, you're very flattering! What it means to me, is to do with not *just* space. That's the first thing to say. I think one of the problems is that when people talk about diffusion and spatialisation, it's as though it's something being *added* to the music, that's not already inherent in the music. To me, whether spatial or any other thing that you apply to the material coming from CD or tape, it should be in the spirit of the music. So space, of course, comes into it, but it is not the only aspect. So, for example, you might argue that one of the limitations of analogue tape – which is where all of this started – was noise. The rather restrictive signal-to-noise ratio of analogue tape without noise reduction is about 60 dB. So, the first thing you need to do is make the quiet bits quieter and the loud bits louder, because the composer will in those days have manually compressed the signal anyway. So the quiet bits are quieter than the loud bits, but they're not necessarily *as* quiet, relative to the loud bits, as they should be, and consequently you need to reinstate that. So that's one thing: you need to reinstate dynamic profiles. Now of course if you think about that, you're going to work off something that already has a relative amplitude curve [draws hypothetical amplitude curve] from 'quieter' to 'louder'. What you are actually going to do is expand it; exaggerate it [draws the same amplitude curve exaggerated].

JM: You mention this in your article in *Organised Sound.*[254] To what extent do you think that specific aspect is still necessary with the increased dynamic range of digital recording media?

JH: It is less necessary, clearly. The thing is, how many of us really use even the theoretical maximum of sixteen bits, which is 96 dB? Probably not a lot. And a lot of composers I know apply compression anyway. I did a talk recently where I used lots of CD examples, snippets. I was editing them

[254] Harrison (1998). "Sound, Space, Sculpture - Some Thoughts on the 'What,' 'How' and 'Why' of Sound Diffusion". *Organised Sound*, 3(2): 117-127.

together and looking at them in the editor thinking, 'Why are Normandeau's levels so much hotter than mine all the time?' The answer is, he's compressed it. A lot of people do that as a matter of course for CD release, learning from the pop industry of course. So that's one thing. The [other] point I was going to make is yes, it *is* still necessary, and the reason it is still necessary is because you can't predict what the audience conditions will be like. So that's why you need to be able to intervene in real-time. That's another reason why I'm not a big fan of completely automated diffusion, which we'll come back to, no doubt, later on. If you accept that, in terms of dynamic profile, what you're doing is exaggerating the implied profile that's already on the medium, then the same would apply to spatial control as well. In other words, if you have something which is going: [makes erratic 'pointillist' textural sound], and zapping around all over the stereo stage, then it seems to me perfectly legitimate to exaggerate that erratic behaviour over a much bigger loudspeaker system. Hence doing this kind of thing [makes short rapid hand movements] and wiggling the faders around in a way that will throw the sound all around the room, in an erratic manner. That seems to be perfectly in-keeping with the musical idea. On the other hand, if you had a big, solid, steady drone going on in a piece, it seems to me that you can't just arbitrarily add random spatialisation to that, unless the piece actually already has random spatialisation within whatever medium its in, like stereo. So if its sitting there, a big fat drone, and you want it to be a 'bigger, fatter drone' in a big space, but you don't go 'wiggly-wiggly-wiggly-wiggly' with it! Because that's not what's on the tape. So that's what I mean about being in-keeping, and being appropriate. The problem is that a lot of people think that diffusion and spatialisation are synonymous, and I don't think they are; I think they're different things. They relate, and very often they are connected, very strongly, in a one-to-one relationship. But very often they're not. In some pieces, it seems to me completely inappropriate to throw it around the space, just because one happens to have a big system, if the music doesn't demand or suggest it.

JM: It seems fairly clear that presentation over one or more loudspeakers is an inescapable fact of electroacoustic music, and that this fact precedes any aesthetic concerns. Which aspects of sound diffusion practice do you see as being applicable to all electroacoustic music, and which do you think might be more directly related to a particular aesthetic standpoint?

JH: Actually, I don't agree with that: I don't think it precedes aesthetic concerns. I think it comes *from* the aesthetic concerns. That's precisely what I've just been trying to say. You would find that there are some means of doing things that are appropriate in some cases – to some pieces, or to some sections of some pieces – which are inappropriate to other sections, of even the same piece, in terms of the kind of movements you would use. For me, that is an aesthetic thing because it actually relates to the way that the music articulates itself. The same goes for dynamics. If there is a big structural moment – a climax if you like – it seems to me that the diffuser ignores that at his or her peril. If you want to play the piece, and *that* is the climax, you have to make *that* climax. You can't just arbitrarily play it a bit quieter: you've got to go with what's there. And that all means an aesthetic understanding, and I think that this aesthetic, actually, is part of the compositional process in a way. You can somehow encode this on the tape, by 'implying' it with the kind of movement that you have.

JM: I see what you mean. Let me rephrase the previous question. On a 'pre-aesthetic' level, if you're going to write electroacoustic music – not necessarily 'acousmatic' music in particular but possibly something you would describe as more 'architectonic' in nature – the fact of the matter is that it will still have to be played...

JH: ...over loudspeakers. I understand what you mean now. I agree, it will have to be played over loudspeakers. In which case, the answer to your question – which aspects of sound diffusion practice do I see as being applicable to all electroacoustic music – I would say probably sensitivity to the musical material, as the first thing. That determines how you diffuse. Because there are certain pieces that I diffuse very lightly. I hardly do anything: slow movements, slow shifts, just 'pointing' certain things at certain places. There are other pieces that I could really 'busk' in a huge way and get away with it, because the material in the piece lends itself to that kind of openness of approach. I think that's the first thing; you have to be sensitive to that. Beyond that, I think it's the dynamic profile, and the spatial profile. Anything that is implied on the tape as 'this-must-be-heard' – this movement from right to left, or this sudden surge of dynamic – then those are the things, it seems to me, that you would exaggerate. You would *ensure* that they are heard. The problem is of course, when you sit in a controlled environment like a studio, you *can* hear them. When you sit in an uncontrolled environment, like your average performance space, surrounded by a hundred other people, you don't necessarily hear them: somebody has to 'play' them for you. That's the key thing. But there are a lot of composers who would disagree with that. They say, 'The movements and levels I want are all on the tape. Just press play!'

JM: This is when you start to get into arguments about wanting to have anechoic halls and standardised speaker rigs.

JH: Yes, very nice ideas, but hardly the most practical. But the problem is, everybody who has an experimental space, including places like SARC at Belfast,[255] just because they get something right in there, does not mean that they then have a definitive version of the piece. They have a version of the piece that works in *that* hall. If you take it to another space, you have to do something else, and I think that's the key thing. When I came over to Sheffield to look at the M2 system, I think its fascinating and the control is fantastic – and I want one – but it still has to be programmed afresh for every space you do a piece in. You can't, I think, just take the same way you did a piece, let's say in Sheffield, you couldn't just 'press play' as it were, and then let the M2 do the diffusion in the CBSO Centre [in Birmingham].

JM: Yes, and that's actually not one of the main aspects of our research.

JH: But a lot of people want standardised arrays. I had a furious email argument with – a very nice guy actually – Jean Piché of the University of

[255] Sonic Arts Research Centre, Queen's University, Belfast. See section 4.15.

Montréal. He wrote, on CEC-Conference[256] about five years ago, that all he needed was four very high quality, high powered, speakers, at a sufficient distance from the audience, so that everybody was sitting in what amounted to the 'sweet spot', and everything would 'work out'. But I just *do not* agree.

JM: So perhaps what you would suggest in that situation is that maybe you would do very little [in terms of diffusion], but you would still have to do *something*?

JH: Well, I don't agree with the premise that you compose on four or eight speakers, or whatever it is, and therefore all you need to do in the concert hall is replicate it. Because the point is, you *can't*. I think that's my basic problem. You *cannot* do it: it doesn't work. It may work in theory, and if the acoustic is sufficiently controlled. But how many halls have you been in that have that kind of controlled acoustic? Not very many.

JM: I find that you always have to do *something*. This is reminding me of one of my own pieces – *Graffiti 2* – which is written in 5.1 formation. So, in that sense I guess you could say the spatialisation is all ostensibly 'in there' [Harrison looks sceptical]. But it does use a very broad dynamic range, and I have found that wherever I've played it, I have had to do certain things, whether it's making the front slightly louder at certain points just to emphasise...

JH: Yes, exactly. And, if you go to the average cinema complex – for example, there's one in town I go to fairly regularly – and basically their surround speakers are just 'echoed' all the way down the sides of the hall. But actually that doesn't really work because basically for most people they're just a side image; that's the net result. For very few people are they actually where they're supposed to be, in terms of the surround.

---

[256] Canadian Electroacoustic Community "CEC Conference". Website available at: http:www.concordia.ca.

JM: So you won't get any particularly effective phantom imaging if the speakers are directly to either side.

JH: Exactly, and also the point is it seems to me that, well, I think this is one of the reasons that the specification for 5.1 as I understand it, is very demarcated in terms of what you're supposed to put on each component in the 5.1 array. The LFE is supposed to have *that* and nothing else, and the centre speaker is supposed to have the dialogue and nothing else, and the frontal stereo is supposed to have the soundtrack, the image, the music, *et cetera*, and very little else. The surrounds kind of vaguely interact with that but mostly for effects. So in *Back to the Future* the car zooms across the screen diagonally, it kind of goes from front-right to left-rear or whatever it does. That's about the extent of it. So they're kept fairly distinct in terms of their functions. But most composition usage doesn't want to keep them that distinct. It wants actually to blur those things. It wants to have something that swirls around the space and then it wants something that's very precise and in a very particular location; those kind of things. And I think that's actually quite difficult to do with 5.1, especially when, very often, the centre speaker will – I don't know if you've noticed – you get a stereo pair of speakers, the same model of speaker – and even Genelec have done this I've noticed – the same model of speaker, ostensibly, has a different number of drive units from the left-right pair. So it can't match the sound. And very often, the surround speakers are a smaller unit altogether than the main stereo. So, actually, it's very difficult to get a fully balanced equal surround, for example, in the average 5.1 array.

JM: Yes, if only simply because there are three speakers in front and only two behind.

JH: Actually, that's an interesting question that we'll probably get on to later. Just put that on the back-burner: 'more in-front than behind'. I will come back to that.

JM: Let's turn to the issue of what Clozier describes as 'diffusion-transmission' and 'diffusion-interpretation.'[257] Do you think that there could be a case for a 'diffusion-transmission', where the idea is essentially to attempt – as far as possible – to recreate what the composer might have intended you to hear in the studio (and I'm thinking perhaps more specifically of multichannel works with 'fixed' panning), and a 'diffusion-interpretation', which is rather more for pieces of music that seem to invite the performer to be more creative in the way that they diffuse it?

JH: If I understand the distinction you're making between those two terms, you're saying that diffusion-transmission is, in a sense, this supposed 'replication of what the composer composed', and diffusion-interpretation is something where somebody comes along and 'plays it differently'. If we can just take an instrumental analogy for a moment, the diffusion-interpretation is equivalent to somebody playing a well-known piece at a tempo that is completely abnormal, for example. Somebody could take liberties with it. Is that the distinction you're making? Or are you saying that there's some pieces that are intended to be diffused: they fall into diffusion-interpretation; and there are some that the composer, particularly using multichannel, has a fixed notion of what should happen, and he wants everybody to hear it like that?

JM: It's all of these senses, but if I can just draw back in a point that we mentioned earlier on. Stemming from the simple fact that all of these pieces of music will have to be played through loudspeakers, there's always going to have to be a 'diffusion' of sorts, even if that's all you understand it to mean. And so it occurs to me that we may need different 'degrees of diffusion' if you like.

JH: Well I think there *are* certain 'degrees of diffusion'. That's in a sense what I've been saying right from the beginning. It's not appropriate to do, necessarily, the same kinds of things in any given circumstance, to a whole

[257] Clozier (1998). "Composition, Diffusion and Interpretation in Electroacoustic Music". In Barrière and Bennett [Eds.] *Composition / Diffusion in Electroacoustic Music*: 235.

bunch of different pieces. I went to a conference in the States, and the guy on the desk 'diffused' the pieces; all the pieces. And what he actually did was, he had choreographed, on the mixer, certain fader movements, and he did them irrespective of what material was coming off the medium. So he just went through the same routine: he just went from these faders, to these faders, from these faders to these faders, to the next, the next, and then he started again. And he did it at the same tempo; he did the same pattern, the same sequence. He was just arbitrarily moving the sound, no matter how fast it was coming, no matter whether it was lots of movement implied or no movement implied, he was moving it round the room arbitrarily, in the same sequence, piece after piece after piece. That is not diffusion: that's ridiculous; that's just nonsense. How did I get onto that? The point about that is that, again, it goes back to this thing about you've got to find what's appropriate for the given piece. If you say, 'Let's take this piece; its stereo so all we need is two loudspeakers,' well I think most people – even in the USA, where this has taken a long time to catch on at all – would actually tend to say, maybe that isn't sufficient: its fine if you're in the sweet-spot but not at the back or the front or the side or whatever. So you need more speakers. Well, then you've got to work out how to deploy certain sections of the piece across the four speakers, or the eight, or six, or twelve, or whatever it is you're using. And then you're into diffusion. So inevitably it grows from a need to reinstate, if you like, in a big space, that which is audible in a small space but would get lost in a big space. So that's the basis. I think the diffusion-transmission thing – in the sense of taking that principle and saying 'Let's assume a big space, let's assume a lot more loudspeakers but now, let's *fix* it' – I can see where they're coming from. But, all I can say is that my attempts at doing it like that have proved that it's virtually impossible to do it. My piece *Streams*, which is eight-channel, was written for the 'BEAST main eight' configuration.[258] The problem, first of all, is that if you take it anywhere where they don't understand what that means, you have a problem. Again, at the same conference in the States, they had a typical eight-array for North America, which is to say quad, plus

[258] See Figure 10 (page 135).

the 'diamond'.[259] So one speaker was smack in front of you, and one directly behind. I couldn't use those so I had to end up using the fold-back speakers, turning them to point at the audience, as the mains. And of course they were the worst speakers in the house, and it showed. So there were real problems. Also – a different space – I played it in a cinema in Edinburgh, and because it was a cinema acoustic it was very dry, and all the sounds stayed 'in the boxes'. There was no 'gluing together' of the different locations by the acoustic, which I have had happen in other places. It kind of requires a little bit of 'mixing in the air'. So what you end up doing is putting in extra speakers, to try and have it more fluid. Before you know it, you're diffusing again. And then you're into a different set of problems, because the problem there is that, if you've got an eight-channel source, and every channel of the eight has a fader, how do you do crossfades? You haven't got enough hands! That's where the M2 comes in, because you *can* do that. Anyway, I'm starting to repeat myself now, so let's move on.

JM: Ideally, what kind of results should be attainable in the process of sound diffusion, and to what extent can these be fully achieved at present? What changes might improve the situation, and what issues may need to be addressed?

JH: I think what one is trying to achieve in terms of results is that the audience hears the piece, in one sense, as close to the original as possible. Taking into account the acoustic of the room, the number of people in the hall, how far away they're sitting, these will all contribute and cause that experience to deteriorate. So what you're trying to do is compensate for that, so that they hear it, in a sense, as clearly as they would sitting in a studio. They can't hear the composer's original intention unless you intervene to help them do it, and that's why you have to understand 'what the composer was about'. That is what should be attainable. But that's an aesthetic; an artistic answer. If you're asking me *technically*, then I would say that there are all kinds of things one might want to do. For example – and this is another take on multichannel sources – you might want to have

[259] See Figure 11 (page 136).

certain things that you can diffuse in one manner, at the same time as other things you can diffuse differently. [This, of course, strongly supports the notion of CASS and CLS groupings.] At the moment, with stereo, if you have a drone, and something else going 'skitty-skitty-skitty' all over the room, you're relying on the fact that you can ' nudge' it just enough to make the high frequency stuff move around, whereas the lower frequency drone *doesn't* appear to move around. That's just a 'trick' because of course we all know its *all* coming out of the different speakers.

JM: You're always diffusing 'the whole thing'.

JH: Exactly. What you could do with multichannel, is put your drone in some tracks and the 'skitty-skitty' in some other tracks, and deal with them independently. Which is exactly what I've been trying to do. In my last piece – *Rock 'n' Roll* – I used a six-channel hexagonal array as a 'surround', and I had two solo speakers – the 'mains' if you like , in the BEAST setup – for 'close up' stuff. What I would like to do, in a bigger diffusion rig than we had in Leicester,[260] is be able to diffuse the 'twos' and the 'sixes' completely independently.

JM: That reminds me of Truax's DM-8 system.[261]

JH: Is that the AudioBox[262] essentially?

JM: Essentially yes, I believe.

JH: I've not used it. I've *heard* it, and I know one or two people who have used it. Again, it's a little bit like the M2: it requires an awful lot of time, because what you really need to do is do a diffusion afresh for each space that you install in. And I'm not sure how flexible DM-8 really is, because my understanding is that – at its largest – it's only eight-in, sixteen-out. The

[260] Annual Sonic Arts Network conference, *Sound Circus*, was held at De Montfort University, Leicester, 11-14/6/04.
[261] See section 4.14 (page 229).
[262] Harmonic Functions (2001). "The AudioBox Disk Playback Matrix Mixer". Website available at: http://www.hfi.com/dm16.htm.

M2 is way beyond that, because you're only limited by the number of physical outputs you can attach to your computer, aren't you?

JM: That's right.

JH: So you're well ahead of the game there, I think!

JM: Well... Nonetheless, stereo is still the dominating standard in electroacoustic music. What do you think are the advantages and limitations of this?

JH: The advantages are it's portable and everybody's got it. The limitations are those we've already discussed [laughs]. If you want to publish, if you want to release stuff commercially, that's what you've got to use. Except now we're *starting* to get a bit of the possibilities of 5.1, DVD-A, and that kind of thing.

JM: DVD-A, which seems to have been 'hovering on the horizon' for about five years now...!

JH: No, don't say that! I'm trying to put out a DVD album.

JM: Well, good luck – I would love to see it happen.

JH: Well, some of the new DVD players, I think, can play the uncompressed audio on the DVD-A and Super Audio CD formats. I saw one the other day on a website that claimed to be able to do all this. That would be good.

JM: As we've already touched on, in stereo we only have two discrete audio channels, yet a good diffusion can give the impression of multiple spatially independent streams. How does this happen? What techniques might we use to achieve this?

JH: I think the reason it happens is psychoacoustics, in the broadest sense. You can 'fool' people. For example, if you have some material going on, and you can establish that material in the ears of the listener as relating to a particular part of the space – a particular set of speakers, a particular array, a particular image, or whatever way you want to think of it – and then new material enters, if you can get the entry of the new material, if you can hit it exactly with introducing some different loudspeakers, on top of the image you've already established, you can persuade people that the original image is staying where you already had it, and the new material is coming in on different speakers. That can be done quite successfully, but you have to be exact. It really needs to be exact in every respect. You need to get the right level, and it helps of course if it's spectrally differentiated and frequency differentiated. There's a piece by Vaggione called *Octour*, which is all textural. It's very busy, and eventually all these layers build up – of course it's the *layers* that enable you to do this – so you can set up something, and there's certain particular moments where you're able to introduce new material that's kind of trumpet-like, and you could *swear* its over there on those other speakers. But in fact it's not: everything's everywhere, because it's stereo! So it's psychoacoustics and that's how you do it. But you *use* what the composer has put there in order to do it. You couldn't just add some other speakers 'willy-nilly' and have that happen. It has to be matched to the appearance of new material which, in turn, suggests 'new entry: new location; new image'. And that's where you get the thing to work; where you marry all these things together. It just can't be done by arbitrary means.

JM: That moves quite nicely on to one of the main things I wanted to talk to you about. I believe you have performed some octaphonic diffusion experiments, which I've heard so much about but have not been able to find very much information on.

JH: That's because we didn't document them. But we should do, you're right. Shall I tell you what our thinking has been so far?

JM: Absolutely.

JH: We have various 'prongs' of our thinking, some of which are still very much just 'our thinking' and haven't happened yet. One thing that *did* happen was at BEAST's twentieth anniversary weekend in Spring 2003: we actually put up eighty loudspeakers. The thinking here goes back to my notion about diffusion. The assumption about diffusion on stereo is that, in a big public space, you can't deliver the image correctly. And the *image* of the piece is what we're trying to get over, in all its totality: its dynamic range, its spatial articulation, its full spectrum; the image of the piece as it fluctuates and develops and morphs throughout itself; we're delivering that image. Now, if you do that in stereo using several stereo pairs of loudspeakers because one stereo pair cannot deliver it alone (hence the multispeaker system), then why on earth would we assume that if you're working in eight channels, for example, you can do it with eight speakers? Does the logic of the first experience not imply that what you *actually* need is multiple arrays of eight speakers, not just one? That was the premise. We were going to take an eight-channel piece, and try and diffuse it over multiple eight-channel arrays.[263] And the way we did that was we set up these multiple arrays – and I can draw some of them for you if that would be helpful – a main array; then we put in some that are very close; then we put some that were on the floor; we had a stand in the middle here with eight tiny speakers pointing out like this; then we had eight on the galleries – this was in the CBSO Centre... I don't know if you've ever noticed this, but if you go to the opera, the stage is not flat, and probably it's not flat in the theatre most of the time either. It's raked, so that people sitting in the house see something that appears to be flat, because they're higher up, but in fact it's a rake. So what we did was – and I did this with a dance project – I had the stage with the dancers on it, moving around. They were being tracked by camera, and the movements were translated into sound in space. So we did it like this: these [two speakers] were on the floor; these [two more speakers] were up a bit; these were higher; and these were higher still. So, if you look at it from the side, it looks like this [draws diagram]. The dancers were moving around in this space, but of course to the audience looking

[263] See Figure 28 (page 208).

from the front, they could hear the sound following the dancers around. But if they'd [the speakers] all been flat, it would have maybe just been doing this a bit [makes level left-to-right movements]. But the way we did it, it was more three-dimensional. So, [returning to the octaphonic diffusion rig] we had some on the floor, some higher ones, higher ones, and then some up on the top gantry. We had some right up in the roof, I don't quite remember where. There's a central gantry across the CBSO, so we had two on the floor, two on the first level, two on the next level, and two in the ceiling, like a Ferris-wheel.

JM: Ah, so tilted upwards through ninety degrees.

JH: Absolutely: a vertical array. And we also had a vertical array on the back wall, like a rose-window in a church. The idea of what we did was very simple: Max/MSP, and an eight-channel sound file assigned to one fader controlling every eight-channel array. That's it; piece of cake. And the other thing, of course, in Max, because you can mix the signals just by attaching two signals to one input, so in fact we had eleven 'virtual arrays' over eighty speakers. So these eight, on the floor, were also part of a 'reversed eight'. They were '1, 2, 3, 4, 5, 6, 7, 8,' but they were also '8, 7, 6, 5, 4, 3, 2, 1'. We sub-mixed all the sub-woofers as well, because we didn't have eight of those, or we were running out of channels, or something. So, eighty-eight virtual channels. We were running the MOTU 24 I/O – we had two of them – and we had a 2408 Mark 3, and we were also using the digital optical outputs to get more channels. That [the eight-channels of audio] was all coming from an 8-to-8, which was feeding the computer optically, and then we also had eight channels out of the [MOTU] 828 which fed this little array in the middle here, because that's where the 828 was. So that's how we did it. And then, just a little Peavey box with eleven faders. Now of course the problem with that is, if you've got certain kinds of movements on the piece – and the piece we did was a piece I threw together the night before, literally, so it was no great shakes, and it had a lot of circular motion in it, which I wouldn't normally do in quite that way. The point was that I wanted to find out whether circular motion, through crossfading of arrays, could be turned into, for example, spirals. So you go from ground, through to main, through to higher, through to roof, and back again. So instead of just going round and round and round, it actually climbed up, around the space, and back down again. Kind of like a 'Slinky' – those metal coils that go down stairs – that kind of effect in sound. So that something could come round, through the Ferris-wheel array and then 'stick' up on the far wall. All that kind of stuff.

JM: You're almost pre-empting my next question. To what extent were the perceptual effects of the octaphonic diffusion as effective as those we've talked about in stereo, where you get this kind of 'phantom separation' of sources?

JH: To some extent, it worked. The basic mechanics of it worked, so that was one thing to find out. The basic practicalities of it were addressed, which is to say, it's a hell of a lot of work because of the number of loudspeakers and so on! The big problem was that we only tested it with one piece, and it wasn't really a proper piece; it was a kind of étude. That is problematic because in a sense the 'database' isn't large enough to draw any conclusions from. We need to do it again, with more pieces.

JM: We're going back to the question, 'Why does stereo dominate?' 'Because it dominates!'

JH: Yes, partly. A lot of people did say to me that they really were surprised at how effective the sound was, but you can't draw much from that. We needed to do more; we needed more objective tests; we needed to try different kinds of material. But I'm still convinced that it could work. However, I have to say that I'm no longer working within a regular array of eight, so now I'm thinking more about this 'two plus six'. So, for every one of these 'eights', you could think of a 'six', but you could also think of additional 'twos', or 'twos' that are part of other 'sixes'! And of course in Max/MSP you just draw cables to connect them: it's not a problem. All you need to do is have the fader controlling the right things, and you could do this. You could have an array of six speakers, in one instance, where also two of them are the stereo. You've got *independent* control over the 'two' and the 'six'. That's what I'm thinking at the moment. Musically and aesthetically I want the control over things that are close and intimate in quality, and things that are more general and 'surround' in quality, and I want to exploit that. You can certainly think of all kinds of ways of doing things: it all depends on how you set up the hall, and how you configure the software. But it's actually relatively easy to do. If we were to get an M2 [Harrison makes a suggestive glance], we would be looking to do this kind of thing, because I think it would lend itself to it very well. Still, though, the big problem is going to be the controller interface: what is controlling what, and how?

JM: Keeping that thought in mind, it's probably true to say that throughout the history of sound diffusion, the mixing desk fader has been the primary physical interface. What do you see as the advantages and disadvantages of this, and to what extent do you think alternative control methods might be successfully employed?

JH: The advantage, first of all, is that it's an interface that practitioners of electroacoustic music tend to know fairly well. Ironically, they know it less well if they've been brought up entirely in the digital domain, because it hasn't been as necessary. Whereas those of us who grew up making pieces with analogue tape, you *had* to use faders. You *had* to use the mixer because it was the only way of mixing stuff together, so you got to know what a fader felt like. So that's why it worked, historically, quite well, and has been the dominant one. The disadvantages are, of course, that it only allows certain kinds of access. If you have *physical* faders, they have to be a certain distance apart. They are, therefore, going to be sitting next to some other fader, but not next to that fader over there at the other end of the desk. So if you happen to want to do something with that fader and this fader, and they're half a metre apart, you're stuffed! You haven't got enough hands. That's where other kinds of options might come in. One option, of course, would be automation, or using presets with 'slews' between them to smooth out the transitions. That, of course, is all possible. To me, you see, what you would need would be some kind of data entry thing to call up the preset, some kind of fader or other variable control to set the slew rate, which can be controlled in real time. Because you won't necessarily want the change always to be at the same rate from one preset to another preset. So you need real-time control over a slew, not a signal.

JM: One of the things we've been looking at for the M2...

JH: When will it be on the Mac? That's what I want to know!

JM: That's in the pipeline as well, actually. With the new version we're working on fully platform independent code, which can be compiled for Windows, Macintosh, and Linux. One of the other things we're working on is a bit like a 'multi-band crossfader'.[264] Fader is perhaps a little misleading though: it *could* be on a fader. You can define points along that fader that equate to specific presets if you like.

JH: Like the GRM Tools 'super slider'?

JM: Very much like that. We've also been talking about implementing alternative methods of control. Specifically at the moment, we're thinking of using joysticks, which of course give you two dimensions to work with at once; maybe also foot pedals, and we've also talked about some things that might be done with buttons.

JH: I would see all of those as being potentially useful. I've always wondered why we don't have a 'gas pedal' for volume, for example. It's such a simple idea. Whatever preset you've got, whatever state everything else is in, you've got an overriding thing that can give a boost to all the volumes on the desk. That would seem to me to be relatively easy to implement. Taking the 'driving' analogy further, the notion of 'shifting gear' or, as it were, pedal-shift, something like that. That's a kind of preset

[264] See section 6.3.1 (page 292).

thing, I guess. How quickly the clutch kicks in, is the 'slew' [laughs]. Do you see what I'm saying? There's all kinds of ways of thinking about it. If you're talking about real-time control, we've tried some things with game pads. The big problem with them is that they're terribly imprecise. There's not a lot of resolution, and even when the thing supposedly flicks back to the centre, if you actually map the numbers, what it's throwing out every time it hits the centre is within a quite large range. It's not very precise. So I think there are a lot of problems there. Joysticks certainly are a good idea. The thing about joysticks, though, is they pre-condition your thinking to assume that what they're going to be controlling is spatial location. Which is fine, if that's what you want to do. And it is entirely possible to flip images and cross them over, all that kind of stuff. And of course they're smooth: they operate at the rate you set. But in a sense what you're looking at with everything you're doing is a range of values from *a* to *b*. A fader is actually a pretty efficient way of doing that, if you think about it.

JM: There's a lot of possibilities. I think the restriction we're looking at is not necessarily the fader in isolation, but the fader as a single attenuator for a single audio signal. That becomes a different issue.

JH: But like I say, if you got a fader to control the speed of a crossfade from one preset to another, for example, that's slightly beyond that paradigm already, isn't it? So there are other options. One of the reasons we got quite excited about the video thing – the dance project I mentioned earlier – was the idea that maybe one could sit in one's studio with a video camera above your hands, and you could actually do things just in space; certain kinds of motions which could be mapped. Really, it's just mapping: data streams that can be mapped however you want. The big problem with all of them is that every single possible configuration, is a 'new instrument,' and an instrument takes time to learn how to play. At least with the good old fader we have a reasonable history of having learned to play it a bit. All these new, fancy things – gloves and motion sensors and so on – they all have their place, I think, but I think it's a question of how much time do we have to be able to get to control them sufficiently to *know* what we're going to do. I think that's the big issue.

JM: One of the things I've observed, relating to that, is that a lot of new diffusion systems and tools seem to propose too many, and too radical, changes. It's almost as if they say, 'This is the old way of doing things, we're going to throw that out the window and present you with something completely new.' This has the effect of alienating people. One of the main things we're trying to do with the M2 is to *expand* on the practice of diffusion, rather than suggest an entirely new paradigm.

JH: I think that's wise, actually. That's very much more helpful for people like me. The other thing that you and I discussed before, *a propos* the M2, was that not only do you need the system to be responsive to the input you're giving it, but you also need it to report back what it's doing, or what state it's in. For example, if you're going from one preset to another preset, you need a sensible and efficient way of getting that information to the person driving. That's always a problem, I think. Obviously one way is motorising the faders; that would be an obvious way to go. So that, when the preset is called – as in any automated desk – then the faders zip to that position, or they move to that position at the rate determined by the crossfade rate. So it recalls the position. But you need motorised faders for that, and of course they have other problems. Every time a motorised fader array is reset, you hear the 'clunk' and 'thump' and all the rest of it. On the other hand, one of the things I hate about MIDI faders is, if you have some other thing that is interacting with it, and you send a message to it saying, 'Set the fader value to *x*,' the next time you touch the fader, it's at the previous value, and it will glitch back to the previous value, and only then will it increment. So that's why the motorised faders, I would say, are possibly the way forward. The trouble is, it also bumps the cost up.

JM: Also, as soon as you get into slightly more complex assignments of 'what's controlling what,' then it's not necessarily always going to be the case that a movement of one control will have an obvious effect on all of the others. That paradigm works better if you're working on a 'one fader one speaker' basis. The problem is, as soon as you start doing things in software, you're increasingly bound continue doing them in software.

JH: And then you've still got the reporting problem.

JM: Then you have to deal with reporting on-screen, really.

JH: Exactly, which takes time to read unless you've got some clever graphics or something...

JM: On a slightly different note, you have previously described 'organic' and 'architectonic,' compositional philosophies[265] I suppose they are. To what extent do you think this dichotomy *predates* the context of electroacoustic music within which you mention it?

JH: Hmm [smiles]. Well that is an interesting question, because I wasn't really thinking about it except *in* the context of electroacoustic music. I suppose my first response might well be to say that probably most music before electroacoustic music was architectonic almost by definition, in the sense of western art music anyway. Obviously there are folk and other musics, which are completely different. The assumption behind the question may be that there is some sort of 'rational' decision-making process going on – to decide that 'this will be this' and 'that will be that' – and the minute you do that, you have some way of deconstructing that into intervallic relationships. The point about intervallic relationships – whether they're pitch-to-pitch or duration-to-duration or location-to-somewhere-else – is that *that* becomes the thing that you're measuring, the 'interval'. The thing that is *at* 'this location' or 'that location,' is secondary to the interval. So what you're doing is you're structuring things by intervals, rather than structuring them by the nature of 'what the thing is' in itself. That's why I call it 'quantitative'. The architectonic is quantitative; it's to do with intervals. Where as 'qualitative' is more to do with 'what this thing *is*' in

[265] Harrison (1998). "Sound, Space, Sculpture - Some Thoughts on the 'What,' 'How' and 'Why' of Sound Diffusion". *Organised Sound*, 3(2)

itself, and then what something else is in itself, but not necessarily with the measurement of the difference between them. Does that make sense [laughs]?

JM: It makes perfect sense, and I think the binarism is a very real one. I got to thinking, before the advent of electroacoustic music and any of the technology that made it a possibility, if you take, say, a piece by Erik Satie, and a fugue by Bach… Actually Bach is probably not the best example. In fact, you have mentioned a piece by Stockhausen, one of the *Klavierstücke*, where you've got extremely precise directions to the performer – you give the anecdote of the nine-note chord with five different dynamic levels – to the point where you really do end up thinking, 'You'd be better to get a machine to do this!' because this really isn't a particularly 'human' process at all. On the other hand, you have something maybe a bit more like Satie where, if you're looking at the score – obviously he was working within that framework – you don't have a tempo indication in beats per minute, you don't have any bar lines, you don't have any key signature. You've got directions in the score such as 'on the tip of the tongue,' and so on. It's almost as if the system itself is being used in a way that actually promotes *interpretation* rather than 'reiteration,' if you see what I mean.

JH: I take the point. Obviously, yes, there are people at different degrees along that 'axis of precision,' if you like. Stockhausen, as you mention, is a case in point. In his own work, one of the reasons that he gives for going into electronic music, was precisely to get the accuracy you're talking about. Of course he also wanted to serialise timbre as well. At some point he said something like, 'Give to electronic music what is electronic music; give to instrumental music what is instrumental music.' In other words, let's exploit the different characteristics. If you want absolute precision, do it in the studio; get a machine to do it, effectively. Whereas if you write instrumental music, then allow the player more freedom, or allow the player to have more of an influence. Don't try to write that kind of music [machine-like] for players. From about that time in Stockhausen's work, you see a divergence. The studio work is the studio work, and obviously it's fixed so it's precise.

Whereas he was moving, from the fifties and certainly through the sixties, towards a freer and freer instrumental mindset, where things were interchangeable and you can do this or that, or miss them out, *et cetera*, and eventually to texts. He would say, 'Play a note in the rhythm of the universe' and so on and so forth; *Aus den sieben Tagen* (1968). Of course the ultimate thing is, if you think about it, what he's also on record as saying is that the principle of serialism is actually to do with mediating between extremes. You take a situation and another, different, situation, which could be perceived as its opposite, and you create between them a set of values. Let's say, 'black and white'. You have different shades of grey in between, and that creates a scale. If you then re-order them, that's a series. So what you can actually do, *he says* – this is his philosophical point – is that what you're actually doing with serialist composition, is mediating between extremes, so you don't see black as an antithesis of white; you see it as a *degree* of white. Therefore what happens is, black and white are no longer in binary opposition, they are 'drawn upward into a higher unity,' as he put it. We all now know he thinks he was born on Sirius, and that's slightly worrying! The fact is, this was his philosophical approach: whatever the situation was, he could unify it by dint of applying a serial process. So if he says to some players in an improvising group, 'Play a vibration in the rhythm of the universe,' and then later, 'Play a vibration in the rhythm of your smallest particle,' they are opposites. Let's face it: what could be smaller than your smallest particle, in all practical senses? What could be bigger than the universe, in all practical senses? And then he says, 'Play all the rhythms you can distinguish between the two.' *It's serial!* It's still serial, even though he's allowing a certain kind of freedom. He was not saying, 'This is an F-sharp; play it at *this* dynamic level,' and so on, but it's still coming from the same serialist mindset. Even *Aus den sieben Tagen* can be traced back to serialism, in my view. Anyway, I picked up on the Stockhausen link because you said, 'Get a machine to do it.' At that stage in the fifties Stockhausen was saying, let's do this in the studios, because we've reached an *impasse* here, as to what you can achieve with real players: they can't actually play five different dynamic levels in a nine-note chord and have them heard, and be accurate; it's just impossible. You can do it in the studio: no problem. So let's now let players do something slightly different. And so he wrote slightly different music.

JM: Of course what players *can* do is put some of their own interpretation into the music. That's why I asked to what extent you think that the binarism between the 'more-or-less organic' and 'more-or-less architectonic' ways of thinking, has existed for a long time in music.

JH: Well it *has*. But the boundaries were moved when you became able to capture real-world sounds that existed beyond the 'normal' frame in which music happened, I think. Because everything prior to that, which existed within the 'normal' bounds within which music happened – which was notatable – therefore conformed to certain kinds of norms, and therefore could be measured against something else that also inhabited that same world. What *breaks* the barrier, is where you can go out and capture the sound of a locomotive (or pick another Schaefferian example) and actually incorporate that into what has previously been 'the musical domain'. That pushes the boundaries suddenly, markedly, away from where they were. So yes, there's always been a range of composers operating, from the very intellectualised to the more improvisatory (say Chopin or Liszt). You can situate composers somewhere along that axis, but until you get beyond having to use either notation and/or instruments designed for the purpose to reproduce the sounds you want – i.e. until you get recording – it seems to me that the boundary can only go so far; it's fixed. The minute you can record stuff that goes beyond that, then, 'Whoosh!' the boundary shoots back about a million miles!

JM: So, in a sense, the advent of recording has opened up the full possibilities of that spectrum between the 'organic' and the 'architectonic'...

JH: I would say so.

JM: ...Whereas before it was always necessarily confined to, usually, 'people making sounds'.

JH: And therefore you had to specify what it was you wanted from them, which meant that you were, of necessity, tending towards an architectonic thing. However improvisatory your aesthetic might be, you still have to write down that you want *this* note played at *this* level, for *this* long, on *this* instrument [laughs]. To go with *this* other note, *this* loud and for *this* long, on *this* other instrument! *Et cetera*. Does that make sense? So, yes, I think it's always been there, it's just that the degree of it is much more marked, this binary thing.

JM: Absolutely. But I think that this is something that's not often acknowledged in the literature. In electroacoustic music it almost seems like the implication is that now we are doing things that have never been done before – which is obviously true, to a certain degree, in that there are things that it is now *possible* to do that it wasn't possible to do before – but perhaps in the way people, historically, have been thinking musically, there may be more similarities than the literature acknowledges.

JH: ...But it is now very much more extreme, I think.

JM: So what implications do these different compositional perspectives have for sound diffusion?

JH: I think that the more architectonic a composer is in his or her mind-set, then the less likely they're going to be overjoyed if somebody comes along and re-spatialises or re-shapes their piece in diffusion. Or as Jean Piché put it, he 'Doesn't want to hear Bobby expressing himself on the faders!' Which is how he characterised diffusion. He wants to hear *the piece*!

JM: Although that might be just as symptomatic of Bobby misunderstanding what diffusion is.

JH: Like our friend in America, doing his little choreographed routine on the faders, yes absolutely.

JM: So that might be an issue as well, that people *don't* have a clear impression of what diffusion should be. It's very easy to think –when you first come across the practice – 'Why do you do *that* if you can have an eight-channel piece with it all in?'

JH: I understand entirely, but I think the answer is, 'To make the music sound better.' That should be the answer: that's *why* you do it! Irrespective of what the format of the piece is. To make the music sound better for more people, in a less-than-perfect listening situation. If you put everybody in a studio, you perhaps wouldn't need to do it! [Laughs]. But you can't do that. It is *because* the venue is considerably larger and less predictable [than the studio] that you need to do it at all.

JM: Might there be a case for distinct organic and architectonic 'styles' of diffusion? What would these involve?

JH: The ideal for the architectonic would be no diffusion. That would be the *ideal*. For the organic, all I would say is that I think that certain musics have such physicality already embodied in them, *embedded* in them, that *not* to continue that process into diffusion seems to be a travesty of the piece. Where you have a real sense of something going [makes a 'whooshing' gesture from front to back], that's such a physical gesture that it's just *screaming* to be exaggerated! So you do it with the faders; you just 'do more.' It's not 'cut and dried,' for me, but I would say that 'architectonic people' don't want diffusion, and everybody else is happy with some or other degree of it. That's the problem!

JM: To what extent do think diffusion measures might need to be taken in the performance of multichannel works?

JH: Of course, the big problem is how many fingers you've got! And therefore, as I was saying, our solution was one fader per eight [loudspeakers] in that particular instance. But there are plenty of other ways; in a flexible software environment like Max, or in the M2, it's not a problem. You can have one fader to control as many channels as you want. You just set up a Group, in your system;[266] in Max you just draw in the patch cords. That's one way you could do it, anyway.

JM: Multichannel media have been widely available for over ten years now, yet stereo continues to dominate. What factors do you believe to be involved in this?

JH: The simple answer to that is the domestic market. That's why stereo dominates. It's commercial: sales. Purely and simply. I'm sure that 5.1 is catching up, particularly in North America and Japan, but it's still not the norm here.

JM: Do you think that there might also be an issue of difficulty in composing in more-than-stereo formats?

JH: Well there's certainly a difficulty in composing in anything that isn't standard. It's very interesting that you and Nikos [Stavropoulos], and others at Sheffield, have done 5.1, and we've all been doing other things like eight, or 'two plus six,' because of course it's impossible to get those things out [released commercially]. If I want my DVD-A out, all my pieces are going to have to be converted somehow; remixed into 5.1, because that's what the format is.

JM: DVD-A will only do 5.1?

JH: Well, let me put it another way: although DVD-A will support eight channels of audio, who has eight channels of audio? *That's* the problem! I looked at the Eventide eight-channel DSP outboard unit: 'Fantastic! Eight

---

[266] See section 5.4.2.2 (page 251).

channels!' Look closely: it *doesn't* output eight channels! It *can*, but it's preset to do things in 5.1 and 7.1. And if you look at the panners in Pro Tools or Digital Performer – where you have a mono signal and you want an output routing that is multichannel – you can do it, but only in 5.1, 6.1, 7.1, *et cetera*. You can't do it in eight equal channels. You can pan between seven of the eight in Pro Tools, but not the 'point-one.' Nuendo is the only one, probably, but even that still has these presets.

JM: Then you have the difficulty – and it's something that you've mentioned in the past – what kind of source material do you use? If you've got mono sources then you'll get amplitude-only panning with no phase differences. If you use stereo source signals, *how* do you pan them?

JH: The way we've been doing that here is that most of us have been busy doing all of our spatialisation in Max, and then just importing it to our sequencer of choice. I do still tend to work with stereo sources, but we're about to buy a Soundfield microphone so we're hoping to be able to incorporate 'real' surround sound images (as it were) as well as the 'constructed' ones.

JM: Barry Truax[267] describes a multichannel diffusion strategy with 'multiple channel inputs where each soundtrack can be kept discrete and projected independently of all the others.' Do you have any experience of such systems? What could be the possibilities of such a system?

JH: That's similar to my 'two plus six' idea; it's very much along those lines. It's just that you may not want to do eight independent tracks. You may want to have a stereo and a six, or you may suddenly want to have two quads. It doesn't really matter. If it's set up in software you can just recall the state for that piece, or every patch for example in – you could have one for doing 'two plus six' and a different one for doing eight-channel. For

[267] Truax (1998). "Composition and Diffusion - Space in Sound in Space". *Organised Sound*, **3**(2)

every piece in the concert you just run a new patch and it automatically configures.

JM: How important (in any kind of diffusion) is image integrity? And what issues might arise when we increase the number of channels on the fixed medium to beyond stereo?

JH: This is really interesting, because it depends what we mean by 'image.' Do we mean by the 'image' of the piece, the image that you retain of the piece as it unfolded in time; the image of the piece that you have in your mind; the *notion* of the piece? Or do you mean the actual 'sonic image'; the 'spatial image?' Because it could be both: they're not necessarily mutually exclusive. But they're not necessarily the same thing. Image integrity seems to be everything, except that as I said earlier, what we're trying to do is give the audience the image of the piece: so image integrity *is* everything. Except that, I think the issue is more is 'the integrity of the image' to be taken to mean there is only one 'version' of it? And that's I think where, again, the architectonic-organic divide might show up: the 'architectonic brigade' would say, 'Yes, there is only one [image] – it's the one I composed,' whereas *I* might be more amenable to be saying, 'Actually, it depends on the hall.' I would prefer this thing to be 'close and frontal,' but maybe for *this*, it's not so crucial where it is, exactly. So you get the *relative* thing: it's a little bit more malleable; it's *organic*; it's *still developing* in the course of performance, just as it's been developing in the course of composition. So I would say that image integrity is *everything*, so long as, for example, if you have hierarchical things of distance, you retain those, but exactly *where* you place them to be 'close' and 'distant' is less important than the *distinction* between them. It's not that [assumes deliberately surprised and irritated tone], 'Hey! This sound came from there, it should have come from here!'; I don't worry about that so much. But I do worry about whether I'm getting the kind of relationships in the piece, globally, that I want. *That* is what I would characterise as 'image integrity' every bit as much as some people might think of it as 'fixed locations' and so on.

JM: In terms of the relative benefits and disadvantages, how does the BEAST system compare to other diffusion systems you might have used?

JH: Like the GRM, like Bourges, like all the French ones really, in BEAST we have lots of *different* loudspeakers. But, of course, if you have a stereo pair, they're a matched pair. So if you're going to have an eight-channel array, they need to be the same, I think, because of the eight-channel image. So what we're looking at doing is we're buying multiple eights [of loudspeakers]. So, if someone's going out to do a stereo gig, fine, they can take two of these, four of those, six of those, two of those, two of those, and a couple of sub-woofers; no problem. On the other hand, if you want to do an eight-channel gig, then you take out eight of these, eight of these, eight of these, sixteen of these, thank you very much! You wouldn't dream of putting a stereo pair on an ATC and a Genelec 1029 – that's not a stereo pair because they aren't matched. So if you've got an eight-channel image, you need eight matched speakers. The big problem, it seems to me, is where you *only* have eight, or where the whole system is made up of the same speakers, but that's not that common.

JM: Many electroacoustic concert programmes are likely to contain works that are highly diverse, both in terms of their technological demands and the aesthetic principles on which they are based. To what extent can current diffusions systems cope with this diversity, and how might they be improved?

JH: Well exactly. That's where software, and presets and recalls come in. The way we've solved it with the [DACS] 3D and BEAST is that you can have preset buttons. When you do the track assignments – the channel assignments for the outputs – the way it works, very simply, is that if you've got something coming into inputs one and two, then inputs one and two can all be assigned to stereo pairs [of outputs] for example, so everything gets it's input from one and two. If you've got eight channels from an eight-channel device, you would put them into eight different inputs and then just assign the eight, and then immediately you put the faders up, that device is running and you get the assignment; it works. That's just like your patching thing where you just hook them up; the matrix page on the M2, in essence.[268]

JM: In your experience (with BEAST and with other diffusion systems), to what extent is the potential diffusion determined by the system itself, and how much of an issue do you think this might be? To what extent do existing diffusion systems embody specific aesthetic approaches to composition and diffusion (perhaps at the expense of others) and how might this be addressed?

JH: It *is* determined by the system itself, by the speakers and so on, but it's mostly determined by the combination of that with the room. My phrase that I've used a few times is, 'You sculpt the sound in the space, but you sculpt the space with the sound.' So it's a reciprocal thing between the sound on the tape, the sound of the system and the particular sound in that space. They're not independent; you can't treat any one of those things as independent of any of the others. A fixed medium piece for eight channels might sound one way on a set of Genelecs in a given space, but on a set of ATCs in a different space, it'll sound completely different. You can't have the same conditions all the time in all respects, and those three things interact: the piece, the system [which includes] the loudspeakers, and the room. They're three *variables*.

JM: I was also trying to come at this particular question from an interfacing point of view. So, in your experience, how much is your diffusion of any given piece restricted – or indeed enhanced – by the fact that you're having to do it on a 'one fader, one loudspeaker' basis, where the faders are arranged, say, [from left to right] mains, wides, *et cetera*?

JH: Well, inevitably it is influenced by that, because there are physical constraints. You can of course re-patch or re-jig and get it to be better, and easier to do certain things. I've done a lot of things in France, for example,

[268] See section 5.4.2.4 (page 254).

where they tend to start at one end of the mixer with the speakers that are the furthest away from you in front, and then [gestures from left to right] the next pair, the next pair, the next pair, and this end [gestures to the right] is at the back. That's absolutely tickety-boo if you want to do lots of front-to-back or back-to-front things, but if you want to get the most significant speakers near the audience to be going [makes antiphonal sounds and hand gestures], you'll find that there's a pair there, a pair there, a pair there, and a pair there [indicates four completely separate areas of an imaginary mixing desk]; you can't get that interaction, which is why we don't do that in BEAST. We put those main speakers in a group on the faders so that you can do it easily. But then of course it means that if you want to go front-to-back, it's awkward. So, nothing's perfect! Unless you could suddenly re-assign the faders mid-diffusion. *You* could do that, couldn't you? [Laughs].

JM: The bottle-neck as I see it, is the fact that most systems work on a 'one fader, one loudspeaker' basis. While you'll have the advantage of knowing exactly what everything does (that's the positive side), the negative side is that you essentially have an arbitrarily organised set of controls that are fixed. So there *will* be things that are just going to be very difficult to do.

JH: That is true.

JM: If you could imagine an 'ideal' diffusion system, what would it be like? What does the future hold for sound diffusion?

JH: An ideal diffusion system would be one that makes the music sound good. And what the future holds for diffusion is that I hope to God we carry on with it!

JM: Any developments you see as somewhat inevitable?

JH: I think the move will be, more and more, to multichannel, and therefore an attempt to 'fix' everything. And people will therefore think that diffusion really isn't necessary after all. On the other hand, even in America now they do a lot more diffusion, so in one sense I don't think it's going to die. I think it's seen as a necessary thing and that there is still some need for live interaction, at some level, in performance.

JM: So you could say it's *as* necessary as having a sound engineer – you need to have some 'on hand' to make sure that it all goes according to plan.

JH: I'm not sure about the 'sound engineer' analogy. Do you mean somebody who does sound at a gig?

JM: Just the fact that you have to have someone there, controlling it, rather than it being 'automatic'.

JH: Yes, yes. I think so. Otherwise you would just go into a room full of loudspeakers and sound would suddenly emanate from them for no particular reason. I think somebody needs to be interacting with that sound in real-time, in that space. And the best place to do that, is in the middle of the audience, or at least as *part* of the audience somewhere, not from some control booth off to the side or at the back. Your ears need to be in the same physical space as the audiences' ears.

JM: Professor Harrison, thank you very much.

MISSING PAGE

Appendix 2: Design of the M2 Sound Diffusion System Control Surface

The M2 Sound Diffusion System Control Surface was briefly described in Chapter 5 and is pictured in Figure 30 (page 241). It was designed and built by the author specifically for the purpose, taking into consideration many of the issues discussed previously. This appendix briefly accounts and documents the design.

A Modular Control Surface

A key objective in the conception of the control surface was to obtain a generic design that would allow performers to select control modules (faders, joysticks, buttons, foot pedals, *et cetera*) appropriate to their needs. Accordingly, the M2 Control Surface is, internally, modular in design. The present prototype consists of two 'M2 Atomic Splitter' modules and eight 'M2 Four-Vertical-Fader' modules. These will be described shortly.

Voltage to MIDI Conversion

As described in Chapter 5, communication between the control surface and *SuperDiffuse* client application is via the MIDI protocol. An Ircam AtoMIC is used to derive the MIDI controller messages required by the client application. This allowed for the development of the control surface prototype to take place relatively quickly, as the AtoMIC was already available at the University of Sheffield Sound Studios.

At present, the *SuperDiffuse* client application expects to receive MIDI CC (continuous controller) numbers 7 and 10 ('volume' and 'pan') on all sixteen MIDI channels as a means to externally controlling each of the thirty-two Physical Faders on-screen. The AtoMIC has been configured such that the control voltages received at each of its thirty-two analogue inputs (reflecting the positions of each of the thirty-two faders) are converted appropriately: analogue inputs 1 to 16 become MIDI CC 7, channels 1 to 16, while analogue inputs 17 to 32 become MIDI CC 10, channels 1 to 16. This AtoMIC configuration is saved as a 'patch' that automatically loads when the AtoMIC is powered on.

Power Supply

The power supplied to the Control Surface is also via the AtoMIC, providing the necessary +5v DC reference voltage that is scaled by each of the thirty-two faders. The two-fold relationship between the Ircam AtoMIC and the Control Surface (i.e. provision of reference voltage and conversion of scaled control voltages to MIDI) is illustrated in Figure 52.

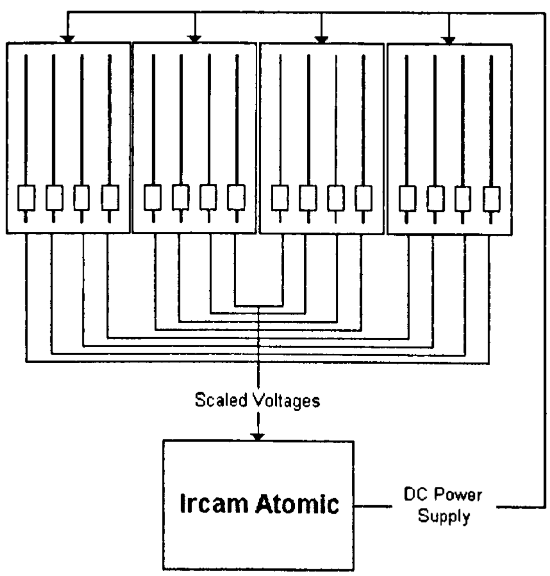

Figure 52. Diagram illustrating the relationship between Ircam AtoMIC and Control Surface modules.

Physical Interfacing

The Ircam AtoMIC has thirty-two analogue inputs, each of which can receive a control voltage (in the range of 0v to +5v DC) for conversion into MIDI CC data. The AtoMIC is also able to provide DC reference voltages, for use if required. The analogue inputs and reference voltages are provided via two 25-pin sub-D connectors, each providing sixteen analogue inputs, reference voltages of +5v, +12v and -12v, and a path to ground. This is illustrated in Figure 53 and Table 5, below.[269]

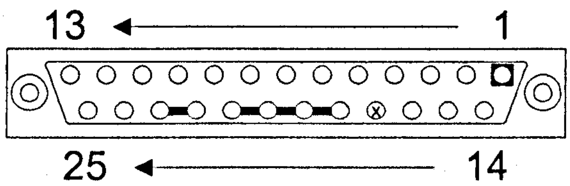

Figure 53. Illustration of the 25-pin sub-D connector, of which the Ircam AtoMIC Pro has two. Pin numberings are shown.

Pin Number	Function
1 – 16	Analog inputs 1 to 16 / 17 to 32
17	Not connected
18 – 21	Ground
22 – 23	+5v
24	+12v
25	-12v

Table 5. Function of each of the pins of the two 25-pin sub-D connectors provided by the Ircam AtoMIC Pro 2.0.

Generic Module Interface

In order to achieve a modular design, it was determined that a generic module interface (GMI) would be useful. This is a low-level physical connection that is able to function regardless of the specific control devices that are connected to it. Effectively, it is an abstraction layer, of sorts. The GMI was devised with the Ircam AtoMIC in mind, but the same physical interface could be used in other scenarios.

It was decided that the GMI should have four control voltage input paths. This would allow four single-axis control devices (e.g. faders) to be hosted by a single module. If dual-axis devices (e.g. potentiometric joysticks) were

[269] Ircam (2002). *AtoMIC Pro Sensors/MIDI Interface User's Manual* (Paris; Ircam): 22.

required, the GMI could host two. The GMI would obviously also need to provide a reference voltage and a path to ground, so six pins in total would be required. For several reasons (economy, availability, increased reliability, scope for future expansion), a 10-pin boxed header was chosen as the physical interfacing method for the GMI. A pin-out diagram for the GMI is given in Figure 54 and Table 6, below.

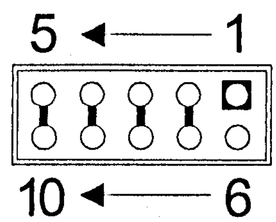

Figure 54. Pin numberings for the 10-pin header used for the M2 Generic Module Interface.

Pin Number	Function
1	CV input 1
2 & 7	CV input 3
3 & 8	CV input 4
4 & 9	Ground
5 & 10	+5v
6	CV input 2

Table 6. Function of each of the pins in the M2 Generic Module Interface.

M2 AtoMIC Splitter (M2AS) Module

The purpose of the M2 AtoMIC Splitter (M2AS) module is to 'convert' between the connectors on the AtoMIC and the M2 Generic Module Interface connectors. This printed circuit board is the abstraction layer that makes the GMI generic, as opposed to an interface exclusively designed for the AtoMIC. Figure 55 is a circuit diagram for the M2AS module; reference to Figure 53 and Figure 54 will be helpful.

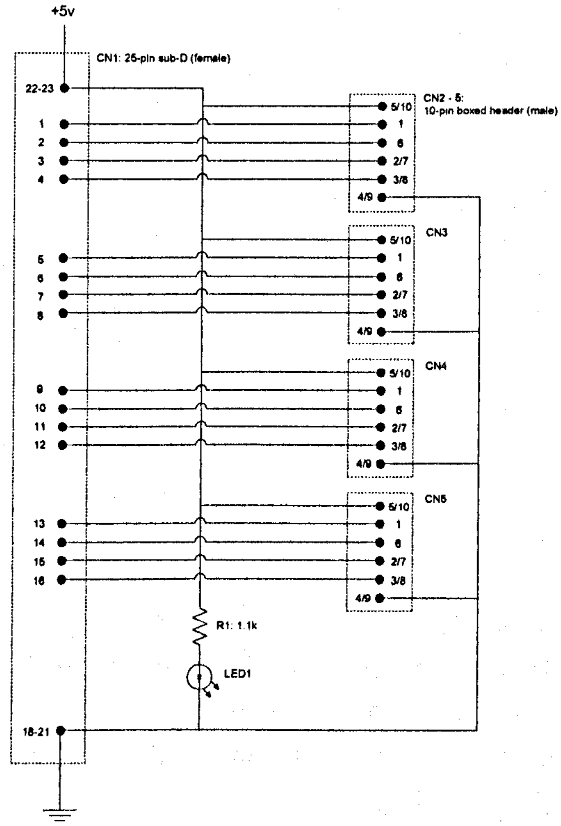

Figure 55. Circuit diagram for M2 AtoMIC Splitter Module (M2AS) v1.0.

It can be seen that the M2AS circuit 'splits' the 25-pin sub-D connector into four Generic Module Interface headers. Because the AtoMIC has two 25-pin connectors, so two M2AS modules are required to access the full thirty-two analogue inputs. An LED is included for the purposes of testing. A facsimile of the M2AS circuit board itself is given in Figure 56; the 25-pin connector (at the bottom-right of the PCB) and four Generic Module Interface connectors (above, arranged horizontally) can clearly be observed.

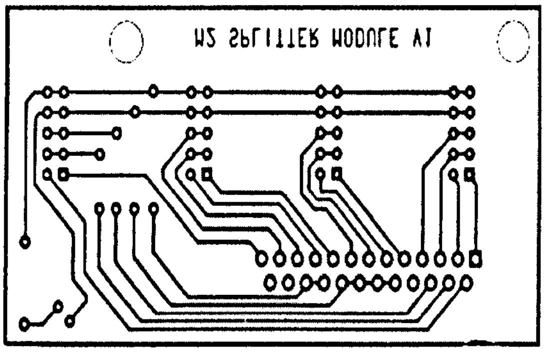

Figure 56. Printed circuit board design for the M2 AtoMIC Splitter module v1.01, as viewed from component side (actual size). Dimensions: 70 x 45 mm.

M2 Four-Vertical-Fader (M2FVF) Module

The M2 Four-Vertical-Fader (M2FVF) module is a printed circuit board that hosts four ALPS 100mm-throw slide potentiometers (faders) aligned vertically; eight such modules are included in the current Control Surface prototype, giving a total of thirty-two faders. The circuit diagram and a facsimile of the PCB itself are shown in Figure 57 and Figure 58, respectively. In Figure 58, the Generic Module Interface can be seen at the top of the PCB, with housings for each of the four faders positioned below.

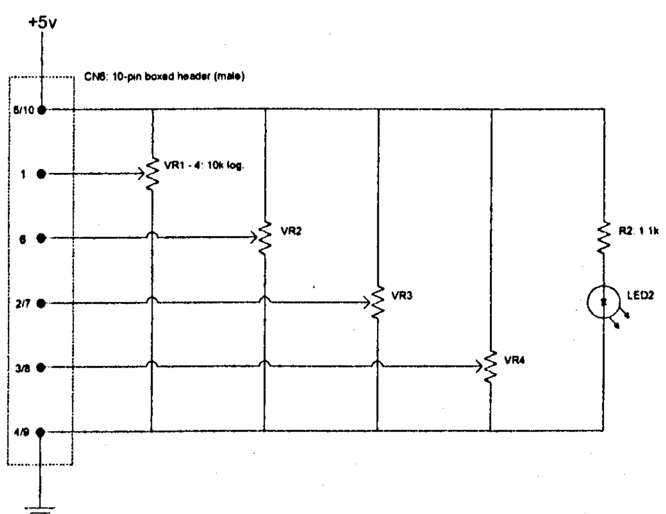

Figure 57. Circuit diagram for M2 Four-Vertical-Fader Module (M2FVF) v1.0.

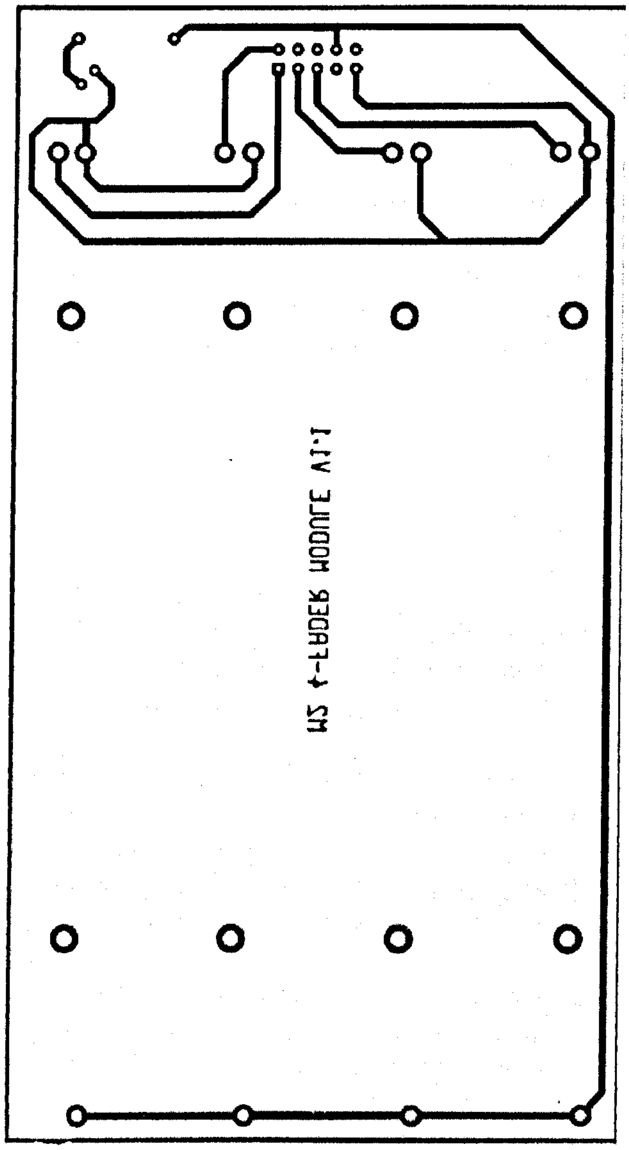

Figure 58. Printed circuit board design for the M2 4-Fader Module v1.1, as viewed from component side (actual size). Dimensions: 80 x 150 mm.

Internal Control Surface Architecture

Figure 59 describes the relationship between the two M2AS and eight M2FVF modules and the Ircam AtoMIC interface. Internally, the individual PCBs are connected with 10- or 25-core ribbon cable. It can clearly be seen that the control modules themselves are *abstracted* from the AtoMIC via the M2AS modules. This means that alternative control paradigms can easily be implemented in place of the vertical faders, so long as the Generic Module Interface is observed. Equally, the AtoMIC itself could be dispensed with in a future design (it is neither the most elegant nor cost-effective solution, but was convenient for the design of a prototype) without necessitating a re-design of the control modules. These are some of the advantages of the modular design.

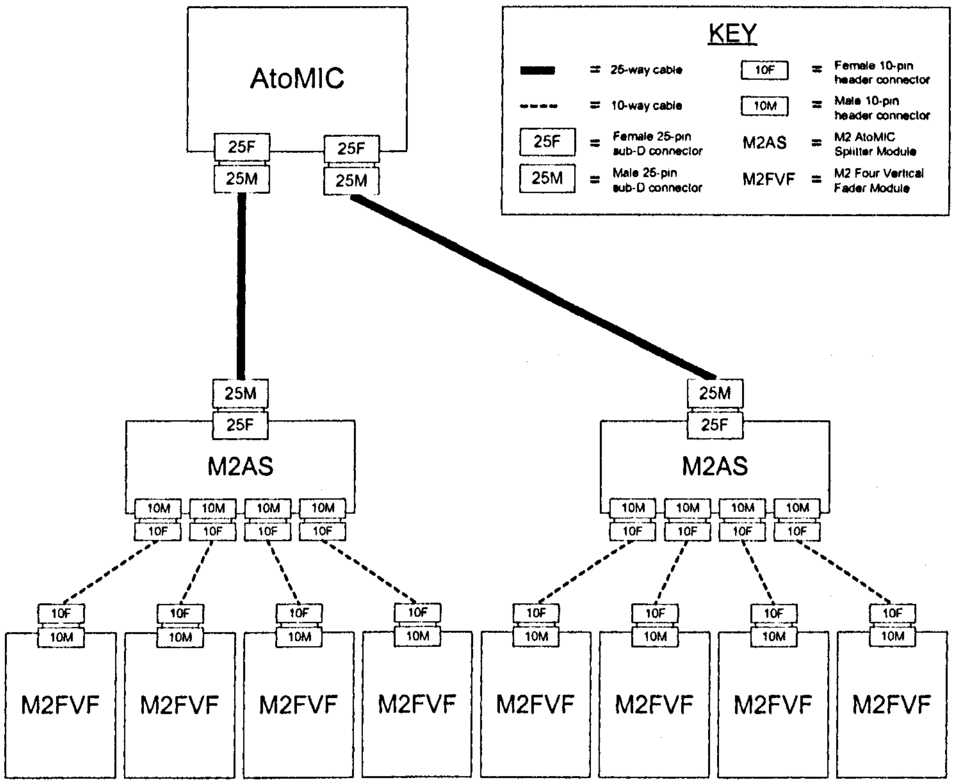

Figure 59. Use of M2AS modules to realise modular hardware architecture using the Ircam AtoMIC voltage-to-MIDI converter, as implemented in M2 Control Surface v1.0.

External Casing

All of the Control Surface components are housed inside an aluminium box designed by the author and built at the Department of Electrical and Electronic Engineering at the University of Sheffield. The box consists of two separate sheets of aluminium, cut to size, and bent to form the top and bottom halves. The top of the box provides thirty-two slots, one for each of the faders – and holes for the two 25-pin sub-D connectors. The bottom of the box simply completes the enclosure. The original CAD designs provided to the workshop are shown in Figure 60 and Figure 61.

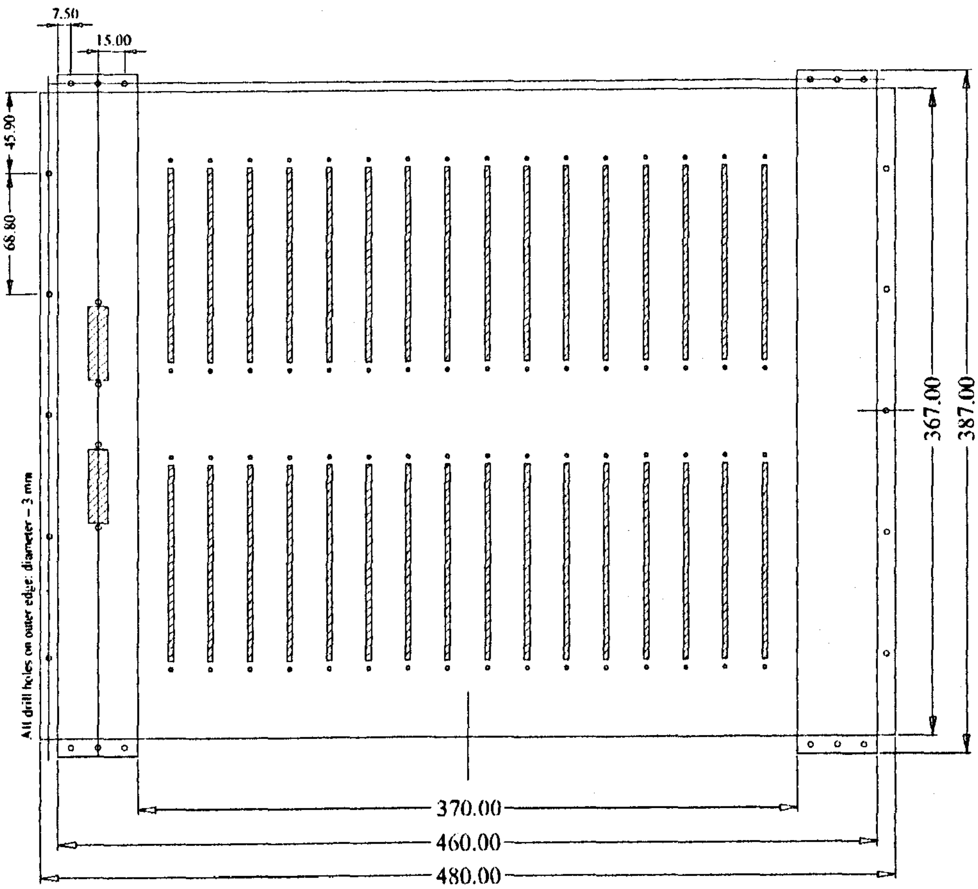

Figure 60. Design for the top half of the M2 Control Surface box. Dimensions in millimetres.

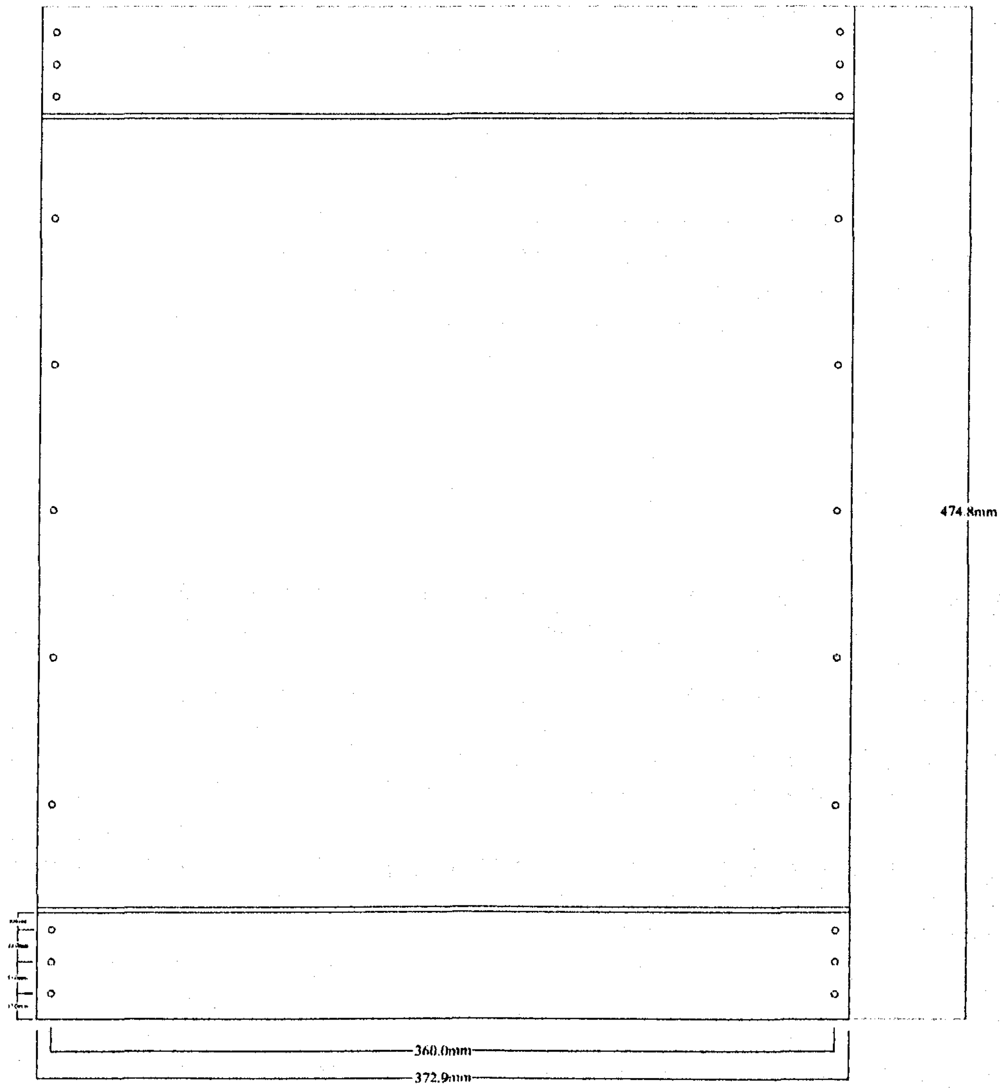

Figure 61. Design for the bottom half of the M2 Control Surface box. Dimensions in millimetres.

PAGE NUMBERING AS ORIGINAL

Bibliography

Austin, L. (2000). "Sound Diffusion in Composition and Performance - An Interview with Denis Smalley". *Computer Music Journal,* **24**(2): 10-21.

Barrett, N. (2002). "Spatio-musical Composition Strategies". *Organised Sound,* **7**(3): 313-323.

Barrière, F. (1998). "Diffusion, the Final Stage of Composition". In Barrière, F. & Bennett, G. [Eds.] *Composition / Diffusion in Electroacoustic Music* (Bourges; Editions Mnemosyne): 204-209.

Berenguer, J. (1998). "Music-space". In Barrière, F. & Bennett, G. [Eds.] *Composition / Diffusion in Electroacoustic Music* (Bourges; Editions Mnemosyne): 219-220.

Blumlein, A. (1931). "Improvements in and relating to Sound-transmission, Sound-recording and Sound-reproducing Systems". Pat. App. No. 34657/31, International Pat. No. 394325. Alan Blumlein Official Website. Available online at: http://www.doramusic.com/patents/394325.htm.

Boesch, R. (1998). "Composition / Diffusion in Electroacoustic Music". In Barrière, F. & Bennett, G. [Eds.] *Composition / Diffusion in Electroacoustic Music* (Bourges; Editions Mnemosyne): 221-225.

Bukvic, I. (2002). "RTMix - Towards a Standardised Interactive Electroacoustic Art Performance Interface". *Organised Sound,* **7**(3): 275-286.

Canadian Electroacoustic Community. "CEC Conference" Mailing List. (Montreal QC.; CEC). Website available at: http:www.concordia.ca.

CDP (2005). "CDP Home Page". Website available at: http://www.bath.ac.uk/~masjpf/CDP/CDP.htm.

Chion, M. (2002). "Musical Research: GRM's Words". (Paris; INA). Website available at: http://www.ina.fr/grm/presentation/mots.light.en.html.

Clozier, C. (1998). "Composition, Diffusion and Interpretation in Electroacoustic Music". In Barrière, F. & Bennett, G. [Eds.] *Composition / Diffusion in Electroacoustic Music* (Bourges; Editions Mnemosyne): 233-265.

Clozier, C. (1998). "Presentation of the Gmebaphone Concept and the Cybernephone Instrument". In Barrière, F. & Bennett, G. [Eds.] *Composition / Diffusion in Electroacoustic Music* (Bourges; Editions Mnemosyne): 266-281.

Cutler, M., Robair, G. & Bean (2000). "The Outer Limits - A Survey of Unconventional Musical Input Devices". *Electronic Musician*. Electronic journal available at: http://emusician.com/mag/emusic_outer_limits/.

DACS (Unknown). "3D Sound Diffusion Console". Website available at: http://www.dacs-audio.com/Custom/3d_console.htm.

Decker, S. (2005). "Shawn Decker - Installations". Website available at: http://www.artic.edu/~sdecker/Finstallations.html.

Dhomont, F. (1995). "Acousmatic Update". *EContact!* 8(2). (CEC). Electronic journal available at: http://cec.concordia.ca/contact/contact82Dhom.html.

Doherty, D. (1998). "Sound Diffusion of Stereo Music Over a Multi-Loudspeaker Sound System: From First Principles Onwards to a Successful Experiment". *SAN Journal of Electroacoustic Music,* **11**: 9-11.

Dolby Laboratories (1999). "Some Guidelines for Producing Music in 5.1 Channel Surround". Website available at: http://www.beussery.com/pdf/buessery.dolby.5.1.pdf.

Dolby Laboratories (1999). "Surround Sound: Past, Present and Future". Website available at: http://www.dolby.com/assets/pdf/tech_library/2_Surround_Sound_Past.Present.pdf.

Eagle, D. (2001). "The aXi0". Website available at: http://www.ucalgary.ca/~eagle/.

Emmerson, S. (1986). "The Relation of Language to Materials". In Emmerson, S. [Ed.] *The Language of Electroacoustic Music* (London; Macmillan Press): 17-39.

Emmerson, S. (1998). "Aural Landscape - Musical Space". *Organised Sound,* 3(2): 135-140.

Emmerson, S. (1998). "Intercultural Diffusion: 'Points of Continuation' (Electroacoustic Music on Tape)". In Barrière, F. & Bennett, G. [Eds.] *Composition / Diffusion in Electroacoustic Music* (Bourges; Editions Mnemosyne): 282-284.

Emmerson, S. (2000). "'Losing Touch?' The Human Performer and Electronics". In Emmerson, S. [Ed.] *Music, Electronic Media and Culture* (Aldershot; Ashgate): 194-216.

Ferrari, L. (1998). *Electronic Works*. (Compact Disc - BV Haast: CD9009).

Field, A. (2000). "Simulation and Reality: The New Sonic Objects". In Emmerson, S. [Ed.] *Music, Electronic Media and Culture* (Aldershot; Ashgate): 36-55.

Gerzon, M. (1974). "Surround-Sound Psychoacoustics". *Wireless World,* **80**(12): 483-486.

Gerzon, M. (1995). "Papers on Ambisonics and Related Topics". Website available at: http://members.tripod.com/martin_leese/Ambisonic/biblio.txt.

GRM (2004). "INA-GRM: Key Dates". Website available at: http://www.ina.fr/grm/presentation/dates.en.html.

Harley, J. (2000). "Sound in Space 2000 Symposium". *Computer Music Journal,* **24**(4).

Harmonic Functions (2001). "The AudioBox Disk Playback Matrix Mixer". (Burnaby BC.; Richmond Audio Design). Website available at: http://www.hfi.com/dm16.htm.

Harrison, J. (1998). "Sound, Space, Sculpture - Some Thoughts on the 'What,' 'How' and 'Why' of Sound Diffusion". *Organised Sound,* **3**(2): 117-127.

Harrison, J. (1999). "Diffusion - Theories and Practices, with Particular Reference to the BEAST System". *EContact!* **2**(4). (CEC). Electronic journal available at: http://cec.concordia.ca/econtact/Diffusion/Beast.htm.

Harrison, J. (1999). "Keynote Address: Imaginary Space - Spaces in the Imagination". *Australasian Computer Music Conference* (Melbourne; 4-5 December 1999). Available at http://cec.concordia.ca/econtact/ACMA/ACMConference.htm.

Harrison, J. (2004). "An Interview with Professor Jonty Harrison". Personal interview with the author, 8th September 2004. See Appendix 1 for full transcription.

Harvey, J. (1986). "The Mirror of Ambiguity". In Emmerson, S. [Ed.] *The Language of Electroacoustic Music* (London; Macmillan Press): 175-190.

Hindemith, P. & Sala, O. (1998). *Elektronische Impressionen.* (Compact Disc - Erdenklang Musikverlag: EK81032).

Howard, D. & Angus, J. (1996). *Acoustics and Psychoacoustics* (Oxford; Focal Press).

Hyde, J. (2004). "Projects - Live Sampling". Website available at: http://www.theperiphery.com/projects/LiveSampling.htm.

Ircam (2002). *AtoMIC Pro Sensors/MIDI Interface User's Manual* (Paris; Ircam).

Kirk, R. & Hunt, A. (1999). *Digital Sound Processing for Music and Multimedia* (Oxford; Focal Press).

Küpper, L. (1998). "Analysis of the Spatial Parameter: Psychoacoustic Measurements in Sound Cupolas". In Barrière, F. & Bennett, G. [Eds.] *Composition / Diffusion in Electroacoustic Music* (Bourges; Editions Mnemosyne): 289-314.

Landy, L. (1999). "Reviewing the Musicology of Electroacoustic Music - A Plea for Greater Triangulation". *Organised Sound,* **4**(1): 61-70.

Lévi-Strauss, C. (1964). *The Raw and the Cooked: Introduction to a Science of Mythology* (London; Pimlico).

Lieder, C. (2001). "Colby Lieder - Instruments". Website available at: http://silvertone.princeton.edu/~colby/Instruments.html.

Malham, D. (1998). "Approaches to Spatialisation". *Organised Sound,* **3**(2): 167-177.

Malham, D. (2001). "Toward Reality Equivalence in Spatial Sound Diffusion". *Computer Music Journal,* **25**(4): 31-38.

Manning, P. (1993). *Electronic & Computer Music* (Oxford; Oxford University Press).

Marinoni, M. (2004). "Marco Marinoni's Home Page". Website available at: http://www.geocities.com/marco_marinoni.

McNabb, M. (1986). "Computer Music: Some Aesthetic Considerations". In Emmerson, S. [Ed.] *The Language of Electroacoustic Music* (London; Macmillan Press): 141-153.

McNabb, M. (1993). *Dreamsong.* (Compact Disc - Wergo Schallplatten: WER 20202).

Messiaen, O. (2000). *Turangalîla Symphony / L'Ascension.* (Compact Disc - Naxos: 8.554478/79).

Mooney, J. (2001). "Ambipan and Ambidec: Towards a Suite of VST Plugins with Graphical User Interface for Positioning Sound Sources within an Ambisonic Sound-Field in Real Time". M.Sc. thesis (Music Technology Research Group; University of York).

Moore, A., Moore, D. & Mooney, J. (2004). "M2 Diffusion - The Live Diffusion of Sound in Space". *Proceedings of the International Computer Music Conference (ICMC) 2004* (1-6 November 2004, Miami FL.): 317-320.

Moore, D. (2004). "Real-Time Sound Spatialization: Software Design and Implementation". Ph.D. thesis (Department of Music; University of Sheffield).

MOTU (2005). "MOTU 24I/O Overview". Website available at: http://www.motu.com/products/pciaudio/24IO/body.html/en.

NTMusic UK (2004). "Magnetic Tape Distribution". Website available at: http://www.ntmusic.org/proj_dmu12.htm.

Obst, M. (1998). "Sound Projection as a Compositional Element: The Function of the Live Electronics in the Chamber Opera 'Solaris'". In Barrière, F. & Bennett, G. [Eds.] *Composition / Diffusion in Electroacoustic Music* (Bourges; Editions Mnemosyne): 321-329.

Parmegiani, B. (1991). *De Natura Sonorum*. (Compact Disc - INA: INAC3001CD).

Parmegiani, B. (1996). *Sonare*. (Compact Disc - INA: E5203-275912).

Puckette, M. (2004). "Software by Miller Puckette". (La Jolla, CA.; University of California, San Diego). Website available at: http://crca.ucsd.edu/~msp/software.html.

Roads, C., Kuchera-Morin, J. & Pope, S. (2001). "The Creatophone Sound Spatialisation Project". (Santa Barbara, CA.; Berkeley, University of California). Website available at: http://www.ccmrc.ucsb.edu/wp/SpatialSnd.2.pdf.

Roads, C., Kuchera-Morin, J. & Pope, S. (2001). "Research on Spatial and Surround Sound at CREATE". (Santa Barbara CA.; Berkeley, University of California). Website available at: http://www.ccmrc.ucsb.edu/wp/SpatialSnd.2.pdf.

Roland-Mieszkowski, M. (1989). "Introduction to Digital Recording Techniques". (Halifax NS.; Digital Recordings). Website available at: http://www.digital-recordings.com/publ/pubrec.html.

Rolfe, C. (1999). "A Practical Guide to Diffusion". *EContact!* **2**(4). (CEC). Electronic journal available at: http://cec.concordia.ca/econtact/Diffusion/pracdiff.htm.

Roth, C. (1997). *Being Happier* (Sheffield; Kingfield Press).

SARC (2005). "Sonic Lab". Website available at: http://www.sarc.qub.ac.uk/main.php?page=building&bID=1.

Satie, E. (1999). *Six Gnossiennes pour Piano* (Paris; Éditions Salabert).

Savouret, A. (1998). "The Natures of Diffusion". In Barrière, F. & Bennett, G. [Eds.] *Composition / Diffusion in Electroacoustic Music* (Bourges; Editions Mnemosyne): 346-354.

Schaeffer, P. (1966). *Traite des Objets Musicaux* (Paris; Seuil).

SIAL (2002). "Sound Diffusion System". Website available at: http://www.sial.rmit.edu.au/Resources/Sound_Diffusion_System.php.

Smalley, D. (1986). "Spectro-Morphology and Structuring Processes". In Emmerson, S. [Ed.] *The Language of Electroacoustic Music* (London; Macmillan Press): 61-93.

Smalley, D. (1996). "The Listening Imagination - Listening in the Electroacoustic Era". *Contemporary Music Review,* **13**(2): 77-107.

Smalley, D. (1997). "Spectromorphology - Explaining Sound Shapes". *Organised Sound,* **2**(2): 107-126.

STEIM (2004). "STEIM Products - LiSa X". (Amsterdam; STEIM). Website available at: http://www.steim.org/steim/lisa.html.

Steinberg (2005). "Steinberg". Website available at: http://www.steinberg.de/Steinberg/defaultb0c4.html.

Stevenson, I. (2002). "Spatialisation, Method and Madness - Learning from Commercial Systems". *Australasian Computer Music Conference* (Melbourne; 6th - 8th July 2002). Available at http://www.iii.rmit.edu.au/sonology/ACMC2002/ACMC2002_Web_Proceedings/103-112_Stevenson.pdf.

Stockhausen, K. (1991). *Stockhausen 3 - Elektronische Musik 1952-1960.* (Compact Disc - Stockhausen-Verlag: 3).

Tarsitani, S. (1999). "Musica Scienza 1999". *Computer Music Journal,* **24**(2): 88-89.

Truax, B. (1998). "Composition and Diffusion - Space in Sound in Space". *Organised Sound,* **3**(2): 141-146.

Tutschku, H. (2002). "On the Interpretation of Multi-Channel Electroacoustic Works on Loudspeaker-Orchestras: Some Thoughts on the GRM-Acousmonium and BEAST". *SAN Journal of Electroacoustic Music,* **14**: 14-16.

Tutschku, H. (2004). "Hans Tutschku". Website available at: http://www.tutschku.com/.

University of Glasgow (2005). "Department of Music: Diffusion System". Website available at: http://www.gla.ac.uk/departments/music/studio/diffusion.html.

Various (1990). *Computer Music Currents 6.* (Compact Disc - Wergo Schallplatten: WER 20262).

Various (2000). *Métamorphoses 2000.* (Compact Disc - Musique & Recherches: MR2000).

Various (2000). *OHM: The Early Gurus of Electronic Music.* (Compact Disc - Ellipsis Arts: CD3670).

Various (2001). *Cultures Électroniques No. 15: 28e Concours International de Musique et d'Art Sonore Électroacoustiques.* (Compact Disc - Mnémosyne Musique Média: LCD278074/75).

Vennonen, K. (1994). "A Practical System for Three-Dimensional Sound Projection". Website available at: http://online.anu.edu.au/ITA/ACAT/Ambisonic/3S.94.abstract.html.

Waters, S. (2000). "Beyond the Acousmatic: Hybrid Tendencies in Electroacoustic Music". In Emmerson, S. [Ed.] *Music, Electronic Media and Culture* (Aldershot; Ashgate): 56-83.

Waters, S. (2000). "The Musical Process in the Age of Digital Intervention". *ARiADA* **1**. (University of East Anglia). Electronic journal available at: http://ariada.uea.ac.uk:16080/ariadatexts/ariada1/content/Musical_Process.pdf.

Westerkamp, H. (1996). *Transformations.* (Compact Disc - Empreintes Digitales: IMED 9631).

Windsor, L. (2000). "Through and Around the Acousmatic: The Interpretation of Electroacoustic Sounds". In Emmerson, S. [Ed.] *Music, Electronic Media and Culture* (Aldershot; Ashgate): 7-35.

Wishart, T. (1986). "Sound Symbols and Landscapes". In Emmerson, S. [Ed.] *The Language of Electroacoustic Music* (London; Macmillan Press): 41-60.

Wishart, T. (1996). *On Sonic Art* (Amsterdam; Harwood Academic).

Worrall, D. (1998). "Space in Sound - Sound of Space". *Organised Sound,* **3**(2): 93-99.

Worrall, D. (2002). "David Worrall Home Page". Website available at: http://www.avatar.com.au/worrall/.

Wyatt, S., Ewing, C., Kilbourne, J. C., Oehlers, P., Pounds, M. & Warde, A. (1999). "Investigative Studies on Sound Diffusion / Projection". *EContact!* **2**(4). (CEC). Electronic journal available at: http://cec.concordia.ca/econtact/Diffusion/Investigative.htm.

Zicarelli, D. (2004). "Max/MSP for Mac and Windows". (San Francisco CA.; Cycling '74). Website available at: http://www.cycling74.com/products/maxmsp.html.