This is a repository copy of *Resound: a design-led approach to the problem of live multi-loudspeaker sound spatialisation*.

White Rose Research Online URL for this paper:
http://eprints.whiterose.ac.uk/80582/

Proceedings Paper:
Mooney, JR and Moore, D Resound: a design-led approach to the problem of live multi-loudspeaker sound spatialisation. In: The Royal Musical Association Conference 2008. Royal Musical Association Conference 2008, 15 - 18 July 2008, Aberdeen, UK. . (Unpublished)

Resound: A Design-led Approach to the Problem of Live Multi-loudspeaker Sound Spatialisation

James Mooney (Newcastle University, Culture Lab)
David Moore (University of Sheffield, Dept. of Music)

This paper is based on a presentation given at the Royal Musical Association conference, Institute for Musical Research, University of Aberdeen, UK on Tuesday 15th July 2008.

Introduction

The process of live sound spatialisation has many unpredictable parameters. Large numbers of loudspeakers with variable frequency responses, hastily configured routings, live and recorded audio sources that change between works, unforeseen technical requirements, compatibility issues, unfamiliar control interfaces, limited setup, rehearsal and change-over time, complex power and signal routing schemes, long cable runs and unpredictable venue acoustics are all issues familiar to performers of electroacoustic music.

Unsurprisingly, the articulation of space in a live, multi-loudspeaker scenario can be problematic. Even 'simple' CD-only diffusion requires dexterity on the part of the performer and considerable technical planning. If we consider multiple sources, with more than two channels of audio, each source to be diffused independently under the control of a single performer, the difficulties are compounded. Now consider that we want to stage a live performance of multiple works where each item on the programme has different requirements, and we have a real problem on our hands. A solution to this problem is needed.

In *Towards and New Architecture*, Corbusier observes that a problem, clearly stated, naturally yields solutions through the process of design. Thus, the aeroplane is the logical conclusion to the problem, clearly stated, of sustaining flight. This ethos has been adopted by the authors in beginning to address the 'problem' of sound spatialisation, and in the ongoing development of Resound, a system comprising a bespoke hardware design and open-source software.

This paper begins by stating the problem of sound spatialisation. Against this background, the Resound system is presented in design and implementation.

Sound Spatialisation

Sound spatialisation, or sound diffusion as it is sometimes called, is the process by which sound from electrical or electronic sources is presented to an audience via loudspeakers, as in concerts of electroacoustic music for example. In presenting stereo music from CD, two loudspeakers is theoretically all we need. However, in a large concert hall the experience of two loudspeakers wouldn't be particularly engaging. So, we place additional pairs of loudspeakers at various locations throughout the venue. This means that we have to have some way of controlling the levels of sound from each loudspeaker, and that's where the sound spatialisation system comes in.

What is a Design-led approach?

> Designers tackle ill-defined problems. Their mode of problem-solving is solution-focused. They [...] translate abstract requirements into concrete objects.[1]

> The airplane shows us that a problem well stated finds its solution. To wish to fly like a bird is to state the problem badly, and Ader's "Bat" never left the ground. To invent a flying machine having in mind nothing alien to pure mechanics, that is to say, to search for a means of suspension in the air and a means of propulsion, was to put the problem properly: in less than ten years the whole world could fly.[2]

Figure 1. Ader's *Eole*, a.k.a. "Bat"[3]

The key point of the first quotation is that we approached the design of this system as a *problem solving* process; as a way of finding a solution to a problem. The point of the second, slightly pompous, quotation is that in order to do that, one must first be clear about what the problem is. In designing the Resound system we began by describing the problems of sound spatialisation in a simple and elementary way, without worrying at this stage about what the eventual artistic or creative outcomes might be.

Stating the Problem(s)

What *is* the problem of sound spatialisation? Or, to put it another way, what are the requirements in setting up and using a multi-loudspeaker system, bearing in mind that we are talking about the whole process from start to finish, not just the creative part?

First, we need to transport all the equipment to the venue, since it's rare to have this kind of system permanently set up. Then, we need to set up, configure and test the system, rehearse, and perform. Afterwards, we need to strike and tidy up all of the equipment. In summary we need to:

- Transport
- Set up
- Rehearse
- Perform
- Strike

---

1 N. Cross (2006). *Designerly Ways of Knowing* (Berlin: Berkenhauser), p29.

2 Le Corbusier (1931). *Towards a New Architecture* (New York: Dover), p113.

3 Image from http://www.ctie.monash.edu.au/hargrave/ader.html

Usually, this all has to be done to a very tight schedule. We have set out to develop the Resound system to make this whole process as quick, easy and efficient as possible. In some ways we have been successful, and in others respects we've come across unexpected problems: Resound is an open-source project, and a work in progress, after all.

The following sections describe some of the problems posed by these requirements, and some of the solutions we've been working on.

Transport

In transporting equipment from place to place there are a number of optimisations that can be made. The images below show the Resound server machine, which is mounted in a single flight case. The flight case has wheels, so although relatively heavy it's easy to move from place to place. Internally the system is shock mounted to minimise the possibility of damage in transit.


Figure 2. Resound server in flight case.

The flight case contains the core of the Resound system, which is a rack-mount server computer running Linux that deals with all the audio processing. The flight case also contains the audio interfaces – there are 24 channels of audio I/O – and a CD player. All of those components are internally hard-wired together, so the flight case is completely integrated.

In summary, the system is:

- Portable
- Integrated
- Robust
- Compact

Set Up

Having arrived at the venue and wheeled the flight case in to place, what issues do we encounter in setting up? Power and signal distribution are the main issues we have to deal with initially. The picture below shows the cabling scenario in the middle of the concert hall at an electroacoustic music concert I recently attended. Because the spatialisation system was based on a standard audio

mixing desk, the spatialisation instrument – i.e. the faders – and the signal routing architecture were one and the same. This meant that all the signal cables had to go to the centre of the hall, which is absolutely the least convenient place they can be if you want to get an audience in there as well.

Figure 3. Tangle of cables typical in sound spatialisation setups.

The image below shows that the inputs and outputs to the Resound system are all on the back panel of the flight case. This can be positioned at the back or side of the room, thus avoiding the tangle of cables at the centre. It can also be seen that the whole flight case is powered from one mains outlet, keeping requirements to a minimum. (We do have designs for a power distribution system for loudspeakers but we haven't found the funding to build it yet.)


Figure 4. Rear view of Resound system flight case.

Spatialisation is controlled from a Resound Client program, which can run on a separate computer

(e.g. a laptop in the middle of the room), with all the audio processing is done by the server. Optionally, a third computer can be used to control the playback of sound files from the server. The computers communicate with each other via UDP using standard Ethernet connections.

Rehearse

Now that everything is set up, the next stage is to rehearse. The main problem here is that we have several different performers to deal with, each with different technical and aesthetic requirements, and as usual we have very little time. Let's imagine our concert programme comprises the following:

- a stereo electroacoustic piece
- a 5.1-channel piece
- an 8-channel octaphonic piece
- a piece for clarinet and live electronics
- an improvisation trio consisting of a turntablist, live sampling and sound spatialisation[4]

So, we have an arbitrary number of live and fixed-medium sources, each with an arbitrary number of constituent source channels (stereo, octaphonic and so on). Time is at a premium so we want to minimise the need for physical re-patching. Ideally we want to hook up all the input sources to the system as quickly as possible and forget about them. With a normal hardware mixing desk this kind of programme could cause all kinds of problems with routing, repatching, and so on.

With the Resound system, of course, all the routing is done in software. Across the whole programme we have four live sources (clarinet, electronics, turntablist, sampling), using a total of 7 input channels; these would be connected to 7 of the inputs of the Resound server and would remain connected for the whole concert. For the stereo, 5.1 and 8-channel pieces we don't need to use any physical inputs: we simply transfer these to the server as interleaved or multiple mono sound files and control their playback from the laptop.

Once we have everything physically connected, we then configure the Resound Client software, save a preset for each programme item, and recall these instantly as and when they are required. Since there is no need for re-patching, the change over time between programme items is minimised.

Perform

So far, we have talked about the practicalities of transporting and setting up, and some of the logistical issues of co-ordinating rehearsal and performance, but we have not addressed the most important question of sound spatialisation: How do you control it?

> A typical performance will focus upon the manual control of sound, often with one fader of the diffusion console controlling one loudspeaker. The speed at which decisions can be implemented and the dexterity of control required [...] clearly influences performance practice. Given the practicalities of very little rehearsal time in often inappropriate venues this basic setup can be either extremely limiting or (in the case of a large system) highly intimidating.[5]

---

4   This last example refers to an ongoing project. See J. Mooney, A. Parkinson & P. Bell (2008). "Sound Spatialisation, Free Improvisation and Ambiguity" (Proceedings of SMC08, Berlin).

5   A. Moore, D. Moore & J. Mooney (2004). "M2 Diffusion: The Live Diffusion of Sound in Space." (Proceedings of ICMC04, Miami).

Or, to put the problem a little more succinctly:

> You haven't got enough hands.[6]

But it's not just that we don't have enough hands; actually we are dealing with a complicated problem, which can be summarised as follows. We take an arbitrary number of audio sources (CD, live instruments, microphones, etc), and assume that each of those sources might have an arbitrary number of channels (stereo, quad, 5.1); likewise we take an arbitrary number of loudspeakers, assume that those loudspeakers are arranged in an arbitrary fashion, and furthermore that this will almost certainly change from venue to venue. Then, we need to use the human body (possibly the hands) to diffuse those audio channels to the loudspeakers. In real time. And we may want to spatialise two or more sources independently. And we want to preserve the image integrity of each source, such that the stereos stay stereo, the 5.1s stay 5.1, and so on. And the whole process must be carried out to an extremely tight schedule.

When you put it like that, it's quite a difficult problem, and it's not surprising that off-the-shelf mixing hardware isn't up to the job. The Resound GUI software is, in essence, a set of tools designed specifically for that job. These tools are:

- The mix matrix
- The collective
- The behaviour
- The summing principle

Mix Matrix

The mix matrix is a mixing architecture consisting of audio inputs, outputs and cross-point attenuators. The image below shows what a 3-in 8-out mix matrix would look like in the Resound Client software, with the inputs in green, the outputs in red and the cross-points in blue. The black node is a master attenuator for the whole matrix. By dragging the matrix nodes in the software one can adjust and thus manually set the value of each attenuator. The configuration shown below would deliver input 3 to output 5.

Figure 5. Screenshot of Resound GUI Matrix tab.

---

6   Jonty Harrison, quoted in J. Mooney (2005). *Sound Diffusion Systems for the Live Performance of Electroacoustic Music* (University of Sheffield, Ph.D. Thesis), p328.

This architecture allows us to mix any input to any output in any magnitude, so for the purposes of multi-input, multi-output sound spatialisation it's ideal, *except*: look at the number of parameters we have to control. We can't use a fader for each parameter, because we don't have enough hands, so what do we do? This is where the Collective tool comes in.

Collective

The Collective is a tool that allows multiple parameters to be grouped together on a single fader (or any other OSC or MIDI controller). Elements can be added to the collective sequentially by clicking on matrix nodes. The OSC address in each element tells me which parameter it is (e.g. /am/att/2/7 refers to audio-matrix attenuator input 2 output 7). Each parameter in the collective has a scaling factor that can be set to any value between -2 and +2. As the fader is raised, the value of the fader is applied to each parameter in the collective, multiplied by the scaling the parameter's scaling factor in each case. The screenshot below shows three Collectives, and their corresponding influence on the mix matrix with the fader at maximum position.
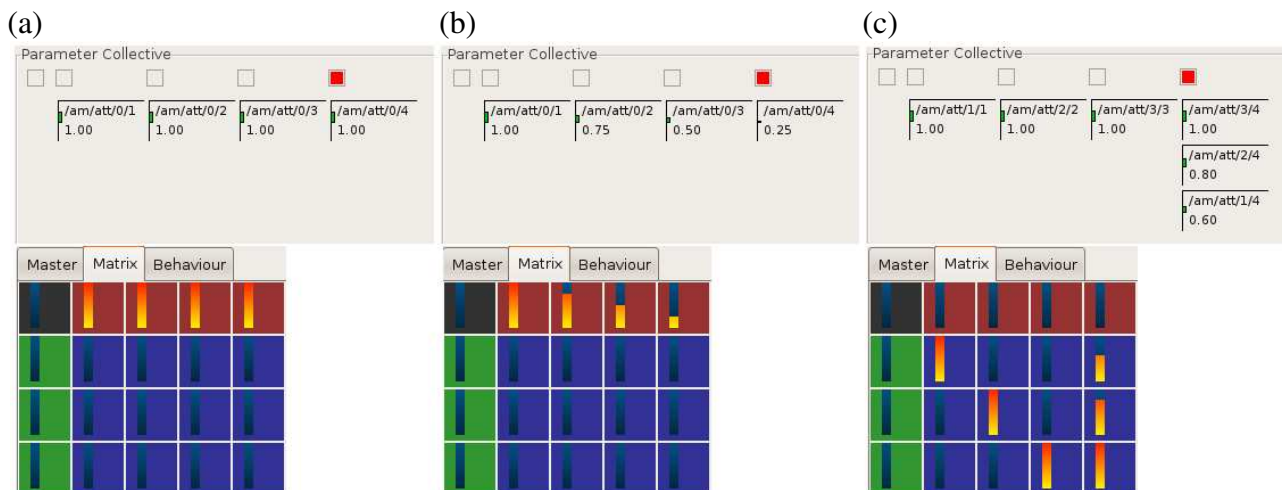

Figure 6. Example Collectives and their corresponding influence on matrix parameters.

In this way the Collective gets round the problem of not having enough fingers, especially with multi-channel sources.

Behaviour

Behaviours allow parameters to be controlled semi-automatically in a variety of algorithmically defined ways. The process is illustrated in Figure 7. First, we create a Behaviour – in this case a Mexican Wave. Then, we assign some parameters to the Behaviour, using the Collective Builder as in Figure 7(b). Notice that the Behaviour itself has parameters, in this case amplitude, frequency and offset (a). These can be set manually or assigned to controls for real-time control (c). The Mexican Wave Behaviour cycles sequentially through the elements of the Collective – this is why the order of elements is important – at the given frequency, as shown in Figure 7(d).

Figure 7. Annotated screenshots demonstrating the Mexican Wave Behaviour.

The Multipoint Crossfade Behaviour is interesting: it is similar to the Mexican Wave, except here we iterate through the parameters of the Collective manually by changing the 'position' parameter (rather than automatically with 'frequency') as shown below. This can be useful for adjusting stereo widths in diffusion, manually fine-tuning the placement of sound sources, and so on.


Figure 8. Annotated screenshots demonstrating the Multipoint Crossfade Behaviour.

Other Behaviours, such as Random, Sinusoidal Wave and Phasor, deal with their Collectives of parameters in different semi-automated ways as suggested by their names.

Summing Principle

In Figure 6 we have two faders – (a) and (b) – whose Collectives act upon the same matrix parameters, so what happens when we use both faders together? The summing principle allows

multiple assignments of matrix parameters to be resolved without conflicts, and as it's name suggests, the influence of each Collective on the mix matrix is summed, per parameter. This is shown in Figure 9, where Faders (a) and (b) correspond to Collectives (a) and (b) from Figure 6. Notice that parameters cannot exceed their maximum value, so it is not possible for this kind of summing to cause signal clipping.
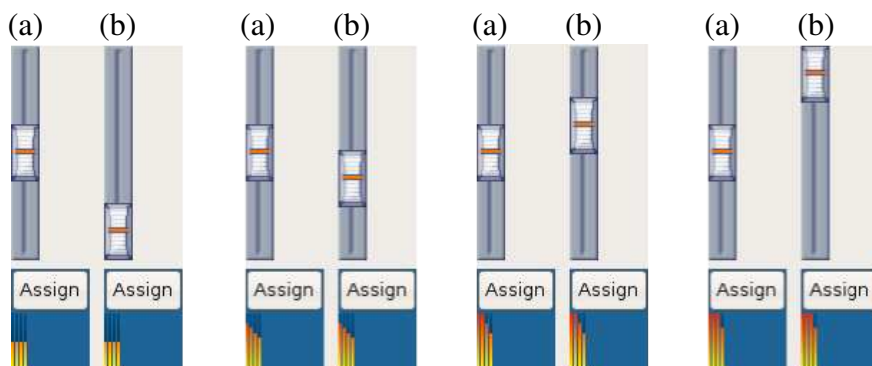

Figure 9. Illustration of the summing principle.
(a) and (b) refer to the Collectives illustrated in Figure 6.

This kind of summing means that Behaviours, which might be quite crude if used in isolation, can be mixed in subtly, and this works quite well in performance.

Another point of note is that if one creates a Collective with negative scaling factors (recall that scaling factors can be between -2 and +2), one can then use that Collective to act subtractively upon the spatialisation, reducing the level of matrix parameters. One can use this, for example, to suddenly shift the image to a single pair of loudspeakers. At the right musical moment this can be dramatic and effective.

Previous, Current and Future Work

As well as doing fairly standard concerts of electroacoustic music there are a number of different directions in which this system is being developed and used. In June 2008, composer Robert van Heumen completed a two week residency with the system, using a joystick to control the Resound client and server using OSC. The buttons on the joystick were used to spatialise different layers of a precomposed piece.[7] We have also staged several concerts using the Resound system as an active instrument in free improvisation, where the ambiguity of not knowing what's coming next gives a different dynamic to the process of spatialisation.[8] In terms of future work, we have begun to look into table-top interfaces, both as a way of controlling spatialisation but also to use multiple audio outputs as part of the interface itself, as a way of auditory cueing.

Obtaining Resound

Resound is an open-source project distributed under the terms of the GNU General Public License. The software, including full source code, can be downloaded from http://resound.sourceforge.net.

---

7   See http://hardhatarea.com/vreemdeling
8   See http://video.google.co.uk/videoplay?docid=4700622507983939