SpaceMaps, Manifolds and a New Interface Paradigm for Spatial Music Performance

Dr. Enda Bates

batesja@tcd.ie www.endabates.net

Trinity College Dublin

BEAST FEaST 2015
University of Birmingham
30 April – 2 May, 2015

Abstract

One of the greatest challenges facing any composer of spatial electroacoustic music is how to adapt their work to different loudspeaker systems, their associated software interfaces, and their implied performance practice. Various multi-channel tools exist which can be adapted for different types of symmetrical arrays, however, these are generally entirely incompatible with the diverse orchestras of loudspeakers associated with the practice of live diffusion. In addition, while there have been numerous attempts to extend or augment the one-fader-to-one-loudspeaker approach to diffusion, developing a system that can flexibly handle the complex routing of many signals in an intuitive and transferrable manner remains a significant challenge.

Manifold-Interface Amplitude Panning or MIAP (pronounced "meeap") is one example of a new design paradigm in which the graphical interface can be arranged without necessarily mirroring the physical layout of the array. MIAP is an expanded implementation of Meyer Sound's SpaceMap spatialization tool for large-scale spatial sound design, developed for the Max MSP environment by Zachary Seldess [1]. While standard panner interfaces can be created using MIAP, so can entirely abstract arrangements, and these can be mapped to arbitrary numbers and configurations of loudspeakers or effects. In addition, the SpaceMap can also be used as a flexible, transferrable configuration and performance tool for live diffusion, in which faders (or other control surfaces) can be mapped to arbitrary arrangements of loudspeakers, much like the concept of the multi-point cross fader previously developed by James Mooney and David Moore for the M2 diffusion system [2]. The SpaceMap could therefore represent a new interface paradigm for the composition and performance of spatial electroacoustic music which is equally applicable to both multichannel and stereo diffusion work, and which could greatly simplify the process of transferring works between different loudspeaker configurations. This paper introduces the MIAP objects for Max MSP through the demonstration of some example diffusion strategies, the multi-point fader, and the transfer of pre-programmed trajectories between different loudspeaker configurations.

Introduction

One of the greatest challenges facing any composer of spatial electroacoustic music is how to adapt their work to different loudspeaker systems, their associated software interfaces and their implied performance practice. The use of multichannel audio and symmetrical arrays of matched loudspeakers, or alternatively the diffusion of a stereo track to a loudspeaker orchestra involve fundamentally different conceptions of spatial electronic music, and how it should be performed [3]. However, even within these two approaches there is still a notable lack of standardization, and adapting works for different venues and performance systems remains a significant challenge. While a number of multichannel spatialization systems have been developed which achieve some level of independence from specific loudspeaker configurations, these remain highly dependent on specific rendering systems [5] and are often limited to symmetrical loudspeaker arrays [4]. Similarly, while the use of stereo sources in diffusion practice offers some level of standardization, the inherently site-specific nature of this approach counteracts this somewhat. The combination of insufficient rehearsal time, and a simple, hardware-based, one-fader-one-loudspeaker approach can also be problematic, as this simple setup can be either too restrictive or intimidatingly complex depending on the size of the array.

From a purely practical level, a typical electroacoustic music concert programme will often involve both types of approach, however facilitating this is a decidedly non-trivial task. While a technique that is independent of loudspeaker layout is desirable, an entirely object-orientated approach is directly contrary to the goals of diffusion [3] and is perhaps somewhat risky given the often less than ideal conditions often encountered during live performances.[1] Given these disparate and conflicting demands, the development of an interface or control method which can satisfy the demands of both stereo diffusion and multi-channel spatialization is a significant challenge.

Summarize requirements for a new system

Any system which can facilitate these disparate criteria will by necessity require flexible and dynamic signal routing, while also separating and abstracting the control interface from the mixing architecture. Historically, the traditional approach to stereo diffusion has consisted of a one fader-to-one loudspeaker approach in which the interface and routing/mixing hardware are fundamentally and physically connected. While relatively intuitive, this approach is also rather inflexible and certain types of movements may be ergonomically difficult to achieve depending on the arrangement of loudspeakers amongst the faders [2]. In recent decades, the use of digital matrix mixers and computers has become commonplace [2] [5] [6] [7]. This type of system can directly mimic a traditional setup using MIDI and OSC control surfaces in place of analogue faders, but with much greater flexibility. Matrix mixers can also readily facilitate other types of routings such as one-to-many fader-to-loudspeaker mappings and dynamic changes in routing or fader assignment. However, this flexibility also greatly increases the complexity of the system. Commonly used matrix-style displays are often cumbersome and far from intuitive, and configuring a matrix mixer to control large numbers of loudspeakers from scratch is a complex task, particularly under the typical time constraints associated with live performances. This issue has been the subject of some research, and the ReSound system (and particularly the concept of the multi-point fader) developed by James Mooney and James Moore is highly relevant in this regard [2]. As Moore points out, a high degree of configurability can be a deterrent to new users if that configuration is overly cumbersome or if not available away from the physical system (in a composer's home studio for example) [2]. The interface needs to be complex enough to support sophisticated performers, but simple enough to be

[1] A controlled environment such as a cinema is perhaps more suited to an object-orientated approach compared to the highly variable acoustics and loudspeaker configurations typically used for musical performances.

learned and used with limited rehearsal time. It should support both physical, real-time diffusion, and also the algorithmic control of pre-programmed trajectories and ideally, the system should also be able to quickly adapt trajectories and movements to different loudspeaker configurations as required.

Finding a solution that satisfies all of these competing demands is a significant challenge. However, one potential solution has in fact been used for over two decades in numerous large-scale, theatrical productions while remaining largely unknown in the electroacoustic community. The SpaceMap spatialization system was originally developed by Meyer Sound Laboratories Inc.[2] for their commercial digital show control system Matrix3 and D-Mitri [8] and was recently ported to Max MSP [9] and Pure Data [10] by Zachery Seldess as a suite of externals entitled Manifold-Interface Amplitude Panning (MIAP) [1]. Created as an authoring and control system for theatrical productions involving very large numbers of loudspeakers, the SpaceMap concept arose out of purely practical considerations such as adapting predefined spatial trajectories to vastly different loudspeaker systems and venues [11]. This is achieved by abstracting the real spatial layout of the loudspeakers onto a two-dimensional topological space known as a manifold, in much the same way as the three dimensions of the planet are mapped onto the two-dimensions of a map [1]. SpaceMaps can therefore be used to directly replicate the physical layout of the array in the same manner as most multi-channel panners, but importantly are not restricted to this. Alternative and entirely abstract arrangements of the loudspeakers are also possible which can potentially be quite useful in the context of elaborate loudspeaker configurations. As the trajectories are stored independently, adapting a particular type of movement to a new loudspeaker configuration simply requires a new SpaceMap, while the trajectories can remain unchanged. In addition, the SpaceMap concept can also be used to create interfaces for the graphical control of matrix mixers, using both pre-programmed trajectories or fader-based, diffusion, and can therefore implement many of the concepts developed by Mooney for the ReSound system, such as the multi-point cross-fader for example [2].

Manifold Interface Amplitude Panning (MIAP)

The following section will provide a brief outline of the fundamental properties of the SpaceMap concept and the associated MIAP externals for Max MSP. For a detailed historical overview of the SpaceMap concept and a more detailed introduction to the MIAP software see [1] and [11]. Fundamentally, the SpaceMap concept consists of an amplitude panning algorithm that uses barycentric coordinates to derive equal power loudspeaker gains among triplet loudspeaker sets. The primary difference between this approach and the well-known VBAP algorithm [12] lies in the abstraction and redefinition of the physical loudspeaker locations onto a two-dimensional, topological space known as a manifold, or a map [1]. SpaceMaps can be created by the user using *Nodes* (several types are available), and *Trisets* (used to link three nodes together and to proportionally distribute a signal between them) [8]. *Loudspeaker nodes* can be positioned on the SpaceMap in a way that mirrors the physical layout of the array, however, arbitrary arrangements are equally viable. In addition, multiple nodes may refer to the same physical loudspeaker. To quote Seldess, a SpaceMap is a "*flexible control surface upon which speaker-speaker relationships are defined without a direct dependency on those speakers' real-work locations*" [1].

Nodes can be joined together to form *Trisets*, within which a constant power panning algorithm is applied (the simplest possible example of such a *Triset* is shown in Figure 1-left). There are four different types of nodes available and *Trisets* can be formed from three of these, namely

[2] Meyer Sound, SpaceMAp, CueStation, Matrix3, and D-Mitri are trademarks of Meyer Sound Laboratories, Incorporated. Meyer Sound, SpaceMap and D-Mitri are registered in the U.S. Patent & Trademark Office.

Speaker nodes, *Virtual nodes* and *Silent nodes*. *Speaker nodes* (shown as a solid black circle) represent a physical output which may be connected directly to a loudspeaker, or indirectly via other effects units and processors. *Silent nodes* (shown as a light grey circle) behave similarly to *Speaker nodes*, except panning a signal to them instead silences the signal. These are generally used to fade in and out signals at the edges of the SpaceMap, as shown in Figure 1-right. *Virtual nodes* (shown as a solid red circle) are connected via weighted links to one or more *Speaker nodes*. A *Virtual node* divides a signal proportionally among its linked *Speaker nodes* and can be used in a few different ways such as creating a panning space inside a group of *Speaker nodes.* This is quite similar to the way many panners simulate sources inside the array by decreasing the directionality of a sound source as it moves toward the centre. *Derived nodes* (shown as a solid green circle) are the only node that do not form *Trisets* and instead simply derive their signal as a linear sum of all the linked *Speaker nodes*. One common usage of *Derived nodes* is to provide a subwoofer feed, as illustrated in Figure 1-right. This example illustrates how SpaceMaps can be used to create traditional panners that directly mimic the physical loudspeaker layouts, in this case standard 5.1.

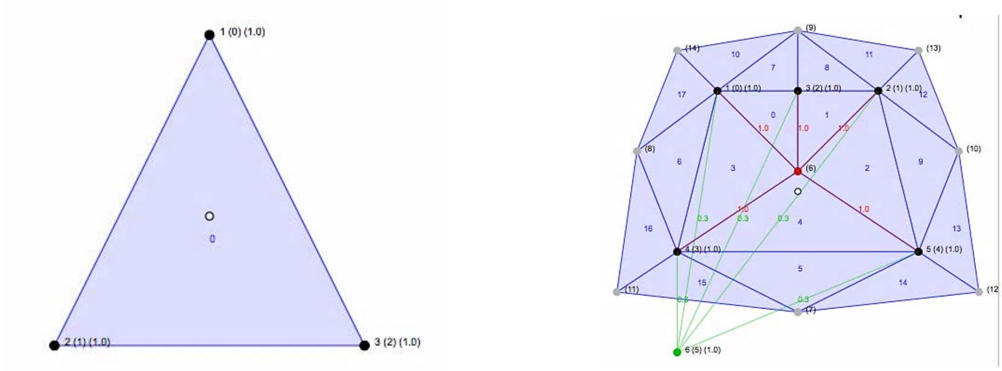

Fig. 1 The simplest possible *triset* (left) and a traditional style 5.1 panner SpaceMap (right)

A source signal can be panned by clicking and dragging in the map using a mouse, or using programmed trajectories which, like the SpaceMaps themselves, are saved externally using the JavaScript Object Notation (JSON) format. As a result, pre-programmed trajectories can be implemented using different arrays, simply by replacing or adjusting the SpaceMap. For example, the 5.1 SpaceMap shown in Figure 1-right could be adapted to include an overhead loudspeaker by simply replacing the central *Virtual node* with a new *Speaker node*. It is also worth remembering that the same loudspeaker can be addressed by multiple nodes. For example a SpaceMap could be created for a simple quad system in which the four inner *Speaker nodes* are routed directly to the array, while the four outer nodes address the exact same loudspeakers, but are first attenuated and routed through a multichannel reverb effect to simulate increasing source distance.

Of course SpaceMaps do not necessarily have to mirror the physical layout of the array, and this perhaps is one of the most powerful aspects of this approach. Figure 2-left illustrates a SpaceMap for a full 11.1 Auro-3D configuration consisting of two, vertically separated five-channel arrays and one central overhead loudspeaker [13]. However, if the user wanted to pan the source signal vertically from the overhead loudspeaker, to all the loudspeakers in the upper array and then finally all the loudspeakers in the lower 5.1 array, this could be achieved by simply panning in a vertical line (shown as a green arrow) between the three *Virtual nodes* of the alternate SpaceMap shown in Figure 2-right. This trajectory could be pre-programmed, however, mapping a physical fader to this vertical line in the SpaceMap is also quite trivial, illustrating one way in which the SpaceMap concept might be used in the context of stereo diffusion.

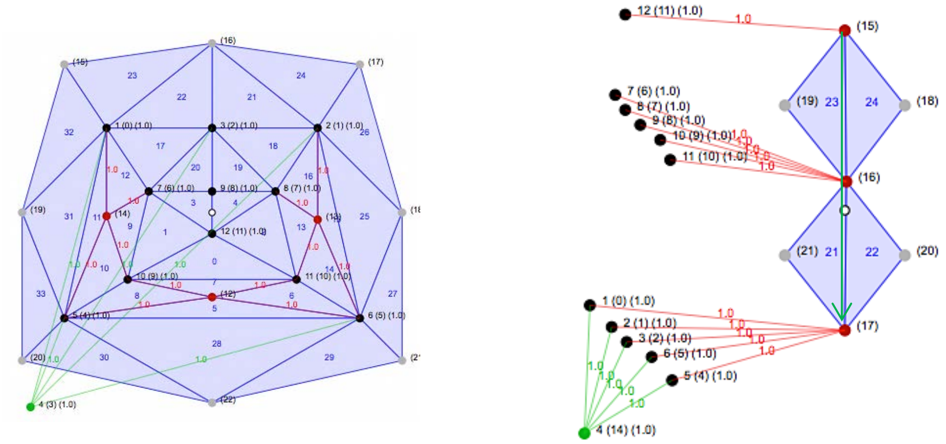

Fig. 2 Alternate SpaceMaps for an 11.1 Auro-3D array

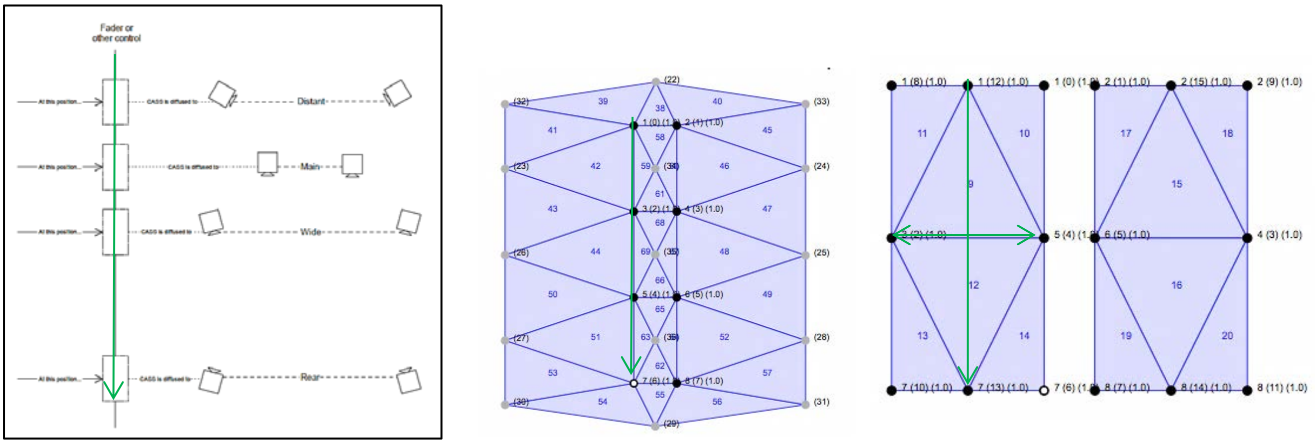

Fig. 3 The multi-point cross-fader [2] (left) and similar SpaceMaps (centre & right)

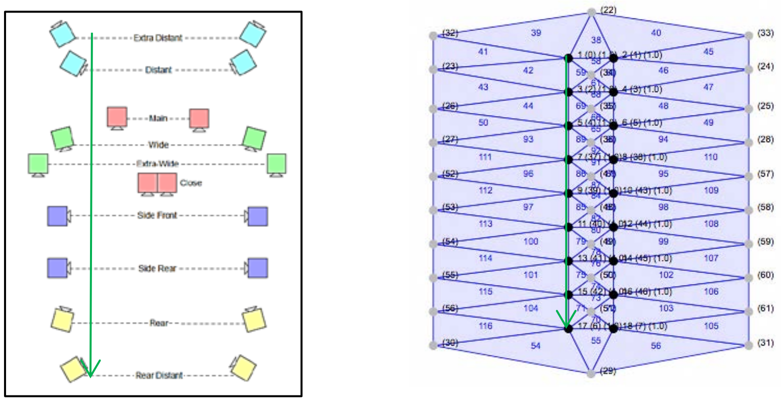

Fig. 4 A front-to-back SpaceMap (right) for a more elaborate system (left) [2]

Diffusion Strategies using MIAP

The alternative SpaceMap shown in Figure 2-right is reminiscent of the multi-point cross-fader concept developed by James Mooney in his work on the ReSound Diffusion system [2], [14] and shown in Figure 3-left. Multi-point cross-faders such as this could potentially be used to enable the diffusion of a work in terms of specific high-level actions, rather than controlling individual loudspeaker amplitudes. This would also allow users to consider loudspeakers in terms of concepts such as 'wide', 'narrow', 'rear', 'mains', etc., rather than addressing venue-specific loudspeaker positions. Potentially, a repertoire of such "diffusion actions" could be developed with a specific

SpaceMap associated with each action. These SpaceMaps could then be modified and adapted to different arrays and venues, but importantly this would not require any change to the stored trajectory, the associated fader movement, or the overall effect of that particular action. Only the SpaceMap itself would change. In this way SpaceMaps could potentially serve as a lingua franca for both diffusion and multichannel spatialization, which is not inherently tied to any one specific type of loudspeaker array.

The following sections will outline some specific examples of such diffusion actions, along with a demonstration of how these SpaceMaps can be adapted for different types of arrays and arrangements of loudspeakers. Note that for stereo diffusion, each SpaceMap would be used twice, with one for each channel of audio. The movement of a stereo source is controlled using the left channel, and this source position is simply linked to the right channel in a mirror image (specifically an inverted mirror image of the horizontal X axis, with the vertical Y axis the same in both). In many cases, *silent nodes* are used as the third node of a *triset*, simply to ensure that the signal will gradually fade out at the edges of the map.

Front-to-Back: This type of movement could be achieved using either of the two SpaceMaps shown in Figure 3 (centre and right) if an external fader is used to move vertically between the four loudspeaker nodes (shown by the green arrow) associated with each channel of the stereo source. The SpaceMap in Fig.3-centre is directly equivalent to Mooney's multi-point fader shown in Fig. 3-left and Figure 4-right illustrates how such a SpaceMap could be modified and adapted to a more sophisticated array, such as the one shown in Figure 4-left. However, this linear pair by pair approach may not be optimal for this type of movement and a simultaneous combination of the front narrow and wide loudspeakers could give better results than a simple sequential distribution. In the SpaceMap shown in Fig.3-right, a vertical movement of the fader (or the programmed trajectory) will again move the source from front to back. However, in this case horizontal movement using a second fader can be used to implement a power-preserving distribution amongst the front-narrow and front-wide loudspeakers as required. The optimal distribution between the front-narrow and front-wide loudspeakers will depend on the actual system in question, however, this can be adjusted by simply altering the horizontal position using an additional, second fader[3].

Unmasking: Unmasking is a commonly used diffusion strategy in which the source is collapsed from a large number or perhaps all loudspeakers, down to just a single pair. This type of diffusion action could be implemented using the SpaceMap shown in Figure 5-left in which an upward vertical movement of the fader (the vertical green arrow) will sequentially and gradually collapse the source from all loudspeakers (bottom position) to just the front-distant pair (top position). *Derived nodes* are used to here to implement a non-power-preserving distribution between the loudspeakers, however a power-preserving distribution could also be achieved using a slightly different map. This example demonstrates how SpaceMaps and one fader can be used to implement a diffusion action which would be quite complex using a traditional one-fader-to-one-loudspeaker mapping.

Sparkling: The use of two faders, or an XY controller, mapped to the vertical and horizontal axes of a SpaceMap is a very efficient way of implementing diffusion actions which are ergonomically challenging using a traditional approach (such as rotations or random distributions for example). Figure 6-right shows one such SpaceMap in which the movement of two faders (whose mapping is again illustrated using green arrows) randomly distributes the source amongst the loudspeakers. In this particular example, stereo cohesion is maintained through the matching of loudspeakers pairs between the left and right sides of the map. However, separate random distributions for the left and right channels could also be implemented by simply altering the *Speaker nodes* on one side. The MIAP externals also include functionality to randomly switch *Speaker node* outputs and this could be used here to achieve a similar effect.

[3] A multi-axis XY controller could also be used in place of two faders in the case of a touchscreen interface.

The same kind of approach could also be adapted to enable the rotation of a source around the array by using a SpaceMap in which the loudspeakers are arranged sequentially in a square. Once again, two faders could be used in sequence to move around the edges of the square in the SpaceMap and so continuously rotate the source around the array, something which is again ergonomically quite difficult to achieve when using a traditional one-loudspeaker-to-one-fader approach.

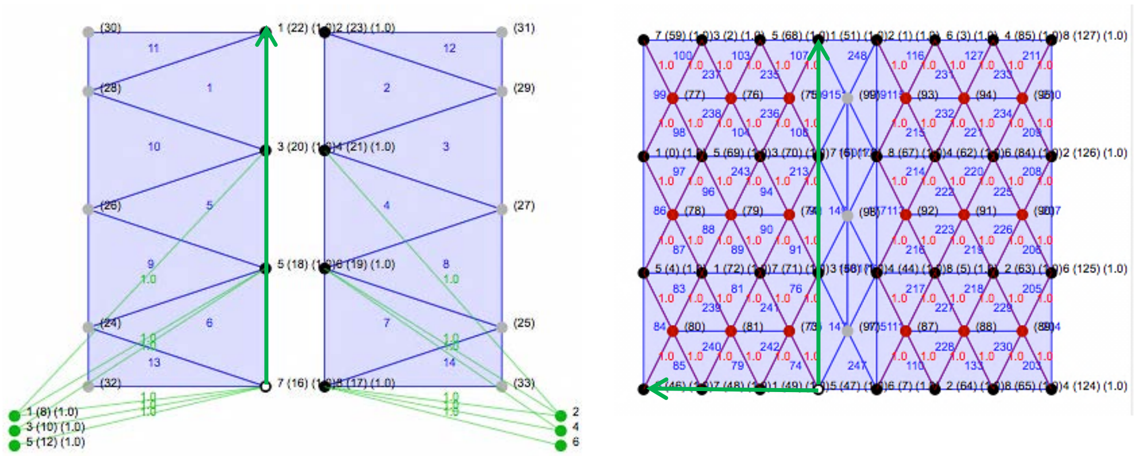

Fig. 5 *Unmasking* (left) & *Sparkling* (right) SpaceMaps for a basic, main-8 diffusion system

Toward a SpaceMap-based Diffusion Performance System

Many existing spatialization systems [2] [4] separate the control interface and signal processing using a client-server based structure. While this can be advantageous for logistical reasons, particularly in a performance context, it was also necessary in the past due to the limited processing power available. A modular approach can overcome these limitations by distributing the signal and interface processing amongst multiple computers; however, as processing power has increased significantly in recent years it is questionable whether this is still necessary. Recent versions of Max MSP support the distribution of patches among different audio threads and therefore can take full advantage of powerful modern multicore processors [9]. In addition, the emerging use of Thunderbolt and Ethernet protocols by soundcard manufacturers has vastly increased the potential number of input/output signals which can be processed in real-time. As such, it is now much more feasible to entirely implement a full diffusion system using open and freely editable patches in familiar and easily extendable and adaptable programming environments such as Max MSP or Pure Data.

To an extent, this approach has already been adopted by systems such as the Zirkonium and BEASTmulch. The ZirkoniumMMK2 software system was developed initially by ZKM for the Klangdom system and later released as a general purpose spatialization program [4]. Loudspeaker configurations are stored by Zirkonium as an XML file (similar to the JSON format used by the MIAP externals) and can be defined and modified by the user, without having to significantly modify previously programmed trajectories. However, while the Zirkonium system is undoubtedly impressive, it is strongly orientated toward symmetrical loudspeaker arrays such as hemi-spherical domes and it is not clear how this system could be used with a diffusion-style loudspeaker orchestra. Zirkonium also defines loudspeaker layouts solely in terms of their actual, physical layout [15] which, in contrast to the SpaceMap concept, limits the extent to which this system can be adapted to non-symmetrical arrays[4].

[4] Interestingly Zirkonium can display periphonic arrays using either a 3D display, or an alternative 2D planer view which is somewhat reminiscent of a SpaceMap manifold.

The BEASTmulch system is another highly developed and fully featured system which was developed as the concert performance system for the BEAST loudspeaker orchestra at the University of Birmingham [5]. Unlike Zirkonium, BEASTmulch can be used with both loudspeaker orchestras and symmetrical arrays and supports both real-time diffusion using a variety of interfaces, and automated trajectories, as well as multichannel techniques such as VBAP and Ambisonics. However, the BEASTmulch system is implemented as a stand-alone application based on a SuperCollider class library which limits the extent to which it can be integrated into a composer's particular compositional workflow. In addition, this dependency is also potentially problematic in terms of maintaining compatibility with operating systems (BEASTmulch is currently only available for Mac OSX 10.4-10.5).

Apart from these issues, both BEASTmulch and Zirkonium are quite complete systems with a number of important additional features that would need to be replicated in any SpaceMap based spatialization system. In a survey of spatialization practices conducted by Peters *et al*, respondents frequently emphasised the importance of three features, namely integration with DAWs, controllability via external controllers, and real-time rendering [16]. While the latter two features are trivial to implement in Max MSP or PD, the integration of MIAP with DAWs requires further development. This is currently achieved in the Zirkonium system using an Open Sound Control (OSC) based plugin [15] and a similar approach could be developed for use with the MIAP externals.

This same survey again highlighted the significant time constraints faced by composers when arranging and optimizing loudspeaker configurations for different venues [16]. The ability to audition and configure the performance system while working away from the venue is therefore another important feature of any new spatialization system. The Zirkonium system includes a binaural renderer which is undoubtedly useful but does not provide any sense of the acoustics of a particular venue. One potential solution to this issue consists of convolving each loudspeaker feed with a binaural impulse response recorded in the venue using that particular loudspeaker. Users could then use the exact same SpaceMap for both loudspeaker and headphone reproduction while getting a sense of that particular acoustic, and reserving precious rehearsal time in the venue to fine-tune physical diffusion actions and/or pre-programmed trajectories. This is quite similar to the virtual loudspeaker approach used in the real-time rendering of spatial audio for gaming [17] and a basic demonstration of this approach is included with the MIAP externals [5].

A number of studies have shown that stereo and octophonic arrays are the most common loudspeaker configurations used by composers and any new spatialization system should therefore support both of these formats [16] [18]. The original MeyerSound system was designed for the panning of multiple mono sources, each of which would be associated with its own specific SpaceMap. Expanding this approach to two-channel stereo is relatively straightforward as demonstrated in this paper, however, the precise way in which this approach can be modified to handle multi-channel sources is less clear. Of course, as the MIAP externals are simply hosted in Max or PD in the usual manner, other multichannel objects (such as Ambisonics externals or the BEASTtools patches for example) could be integrated into a composer's patch without any difficultly. Indeed, maintaining a composition in multiple stems in this way is a useful strategy in terms of adapting a work to different venues and loudspeaker arrangements [5] [19].

Conclusion

The preceding examples have illustrated the potential of the SpaceMap concept to act as a flexible means of codifying both spatialization and diffusion strategies in an intuitive and highly configurable manner. The ability to abstract both the physical layout of the loudspeakers and source trajectories is of particular importance, as it is this process of abstraction that allows existing SpaceMaps to be modified to reflect changes in the array configuration (indeed this type of flexibility and ease of

[5] http://www.zacharyseldess.com/miap/downloads/

modification was one of the primary motivations for the development of the SpaceMap system in the first place [11]).

Spatial audio panning tools frequently use a graphical interface which mirrors the physical layout of the loudspeaker array which, although easy to understand, can also result in serious misconceptions about the nature of spatial audio. When a sound is stereophonically panned to the centre of a 5.1 loudspeaker array for example, the interface visually suggests that the sound will be localized by the listener(s) inside the array, however this is not really the case. In reality, the source is simply reproduced equally by each loudspeaker and so localization will be strongly influenced by several factors, such as listener location. This misconception and the common erroneous equating of graphical layouts with the actual behaviour of spatial audio is a direct consequence of this type of interface design. In effect, interfaces such as these explicitly suggest real space (through the direct replication of the physical loudspeaker layout), while the real and more complex behaviour of spatial audio is only implied. The abstract manifold of a SpaceMap is in contrast an explicitly abstract interface, and so can perhaps avoid some of the pitfalls of overtly graphical notions of spatial sound and of course, spatial music.

One of the most notable aspects of the SpaceMap approach is that it can be used for both stereo diffusion to a loudspeaker orchestra, and for the creation of pre-programmed trajectories using symmetrical loudspeaker arrays. In addition, live diffusion actions, the live control over pre-programmed trajectories (controlling the velocity of precomposed rotations for example), and entirely pre-composed trajectories can all be created and implemented using the exact same interface. This ability to cater for both 'top-down', organic approaches to diffusion, which emphasise the real perceptual effect in the performance space, and also 'bottom-up', 'architectronic' approaches in which trajectories and movements are entirely pre-programmed, is extremely useful and suggests that the SpaceMap has significant potential to act as a unifying format for these quite different conceptions of spatial music.

One of the difficulties in extending the one-fader-to-one-loudspeaker approach to diffusion is the greatly increased complexity in signal routing and interface configuration [2]. A graphical representation such as a SpaceMap is potentially therefore very useful, as it provides a clear, visual depiction of the current signal routing that is far more intuitive than traditional matrix style displays. Even more than this, the SpaceMap could be used as a form of notation and as a means of codifying the practice of live diffusion in terms of specific actions such as collapsing, front-back movements, sparkling, etc. Importantly, these actions and associated SpaceMaps could be tailored to specific loudspeaker orchestras as needed without requiring any change to the physical actions and trajectories created by the composer. This would allow composers to formally associate explicit spatialization strategies with a particular work, and also to rehearse and audition this approach away from the actual venue. This could be greatly facilitated by venues and the operators of loudspeaker orchestras through the provision of SpaceMaps for common diffusion actions that have been fine-tuned to that particular venue and array. The provision of binaural impulse responses, recorded in the venue for each loudspeaker in the array, would be similarly beneficial, as this would enable composers to audition their SpaceMaps, and their associated diffusion actions away from the venue itself.

While many different spatialization systems have been developed in recent years, few of these offer the same degree of flexibility and adaptability as the SpaceMap. The most important aspect of this approach is unquestionably the abstraction of the physical loudspeaker configuration onto a two-dimensional manifold. Ultimately, it is this abstract representation, rather than a direct replication of the physical loudspeaker layout, which enables this highly advantageous adaptability. While the development of a complete spatialization/diffusion system in Max MSP or PD using MIAP is certainly possible (and is currently being investigated by the author), this concept could also potentially be incorporated into existing, well developed spatialization systems. Zirkonium is, for example, a very complete system with lots of important features such as integration with DAWs, the use of Ambisonics, and support for both real-time diffusion and pre-programmed trajectories [4]. However, the interface and underlying spatialization method used by the current version of the Zirkonium is fundamentally based upon the direct, concrete representation of the actual room and loudspeaker configuration [4], which limits its applicability for non-symmetrical arrays and the loudspeaker orchestras associated with the practice of diffusion[6]. The potentially far more abstract SpaceMap neatly avoids this hardwired connection to a specific type of array, and hence to a particular performance practice, and could therefore be a solution to the inherent difficulties associated with adapting works of spatial music to different venues and loudspeaker configurations.

Examples

The SpaceMaps discussed in this paper, along with demonstration patches for Max MSP (OSX only) can be downloaded from the author's homepage at www.endabates.net/EndaBates-Academic.html.

[6] Interestingly a predecessor of Zirkonium, the Topoph system developed by Sabine Schäfer and Sukandar Kartadinata in 1991, used a path-based approach which was not bound to spherical loudspeaker configurations and like MIAP, could flexibly route inputs to an arbitrary number of outputs [20]

Bibliography

[1] Seldess, Z. 2014. "MIAP: Manifold-Interface Amplitude Panning in Max/MSP and Pure Data". 137th Convention of the Audio Engineering Society, October 9–12, 2014, Los Angeles, USA.

[2] Mooney, JR. 2005. "Sound Diffusion Systems for the Live Performance of Electroacoustic Music". University of Sheffield, Department of Music.

[3] Harrison, J. 1998. "Sound, Space, Sculpture - Some Thoughts on the 'What,' 'How' and 'Why' of Sound Diffusion". Organised Sound, 3(2): 117-127, Cambridge University Press New York, NY, USA.

[4] Ramakrishnan, C. 2009. "Zirkonium: Non-invasive software for sound spatialisation". Organised Sound, 14, pp 268-276.

[5] Wilson, S., Harrison, J. 2010. "Rethinking the beast: Recent developments in multichannel composition at birmingham electroacoustic sound theatre". Organised Sound, 15(3): 239-250, Cambridge University Press New York, NY, USA

[6] Clozier, C. 2001. "The Gmebaphone Concept and the Cybernephone Instrument". Computer Music Journal, Volume 25, Number 4, Winter 2001, pp. 81-90, MIT Press.

[7] Beck, S. D., Willkie, B., Patrick, J. 2007. "ICAST: Trials and Tribulations of Deploying Large Scale Computer-Controlled Speaker Arrays". Proceedings of the International Computer Music, Copenhagen, Denmark.

[8] Meyer Sound Laboratories Inc., “CueStation 5.2.0 User Guide”, pp. 115–141 (2011).

[9] Zicarelli et al., “Max visual programming environment”, http://cycling74.com/products/max (accessed March 27, 2015).

[10] Puckette, M. “Pure Data”, Proceedings of the 1996 International Computer Music Conference, San Francisco, USA pp. 224–227 (1996)

[11] Ellison, S. 2013. “SpaceMap: 20 Years of Audio Origami”, Lighting & Sound America, pp. 80–88.

[12] Pulkki, V. 2001. "Spatial Sound Generation and Perception by Amplitude Panning Techniques". ScD dissertation, Sibelius Academy, Helsinki.

[13] Auro Technologies Inc., "Auro-3D", http://www.auro-3d.com/system/concept (accessed March 27, 2015).

[14] Stefani. E. and Mooney, J. 2009. "Spatial Composition in the multi-channel domain: aesthetics and techniques". Proceedings of the International Computer Music Conference, 16-19 Aug 2009, Montreal, Canada. International Computer Music Association.

[15] Wagner, D. 2015. “Manual: Zirkonium MK2 – An Open Source System for Sound Spatialization”, ZKM Center for Arts and Media.

[16] – Peters et al, 2011. “Current Technologies and Compositional Practices for Spatialization: A Qualitative and Quantitative Analysis”, Computer Music Journal, 35:1, pp. 10–27, Spring 2011, MIT Press.

[17] Masterson, C., Kearney, G., Gorzel, M., Rice, H., Boland, F., 2010 “Optimised Virtual Loudspeaker Reproduction”, Irish Signals and Systems Conference, UCC, Cork, June 23–24 2010.

[18] Otondo, F. 2008 “Contemporary trends in the use of space in electroacoustic music”, Organised Sound 13(1): 77–81 _ 2008, Cambridge University Press.

[19] Lyon, E. 2008 “Spatial Orchestration”, Proceedings of the 5th Sound and Music Computing Conference. Berlin, Germany: Universitatsverlag der TU Berlin ¨ (pages unnumbered).

[20] Brümmer, L., Dipper, G., Wagner, D., Stenschke, H. Otto, J. A. 2014 “New developments for spatial music in the context of the ZKM Klangdom: A review of technologies and recent productions”, Divergence Press, Journal, Issue 3, December 2014.