Granular Spatialisation, a new method for sound diffusion in high-density arrays of speakers and its application at the Spatial Audio Workshops residency at Virginia Tech (August 2015) for the composition of the acousmatic piece *Spatial Grains - Soundscape No 1,* for 138 speakers

Javier Alejandro Garavaglia
Associate Professor
London Metropolitan University, UK
jag@jagbh.demon.co.uk

ABSTRACT

This paper was originally submitted for the ICMC 2016 together with the acousmatic piece "Spatial Grains, Soundscape No "1, describing in theory and practice the usage of Granular Spatialisation. Granular Spatialisation is a new and particular case of an on-going development of diverse systems for automatic, adjustable and time-dynamic spatialisation of sound in real time for high-density speaker arrays, and can be therefore contextualised as a further and special case of development in the practice-based research of the author of this paper about the main topic of full automation of live-electronics processes, as explained in [1] and [2]. The paper considers both the theoretical background and the initial phases of practice-based research and experimentation with prototypes programmed to diffuse sound using spatialised granulation. The second part of the paper refers to a recent experience during a residency at the Cube, Virginia Tech, using Granular Spatialisation within an array of 134 + 4 loudspeakers for the diffusion of the acousmatic composition jointly submitted herewith. Seeing that in the past 40 years, the number of speakers for the diffusion of acousmatic music has constantly increased, this paper finds pertinent the main question of this ICMC: "Is the sky the limit?" with regard to the number of loudspeakers that can be used in acousmatic sound diffusion.

1. INTRODUCTION

Based on the general concepts of granular synthesis developed by Truax [3] and Roads [4], which are respectively predicated on Gabor [5] and Xenakis [6], Granular Spatialisation (GS hereafter) transfers the common parameters of a grain (such as grain time, window/envelope type, inter-grain time and grain overlapping time) to the movement *in real time* among loudspeakers within a multi-channel environment.

Although GS works appropriately from a 4.0 surround system onwards, the ideal speaker configurations are those with a large number of loudspeakers located in different areas within a performance space, which can be, for example, a concert hall or a gallery space, the latter mainly for sonic or audiovisual installations.

2. GRANULAR SPATIALISATION: MAIN FEATURES

GS has been conceived to work in both performance and/or installation environments, with at least a 4.0 (quadrophonic) sound output. The directionality of the sound can be in either surround or in any other type of directional arrangement. However, the main goal is to work in loudspeaker configurations, which are significantly larger, for example, the *Klangdom* at ZKM (*Zentrum fur Kunst und Medientechnologie*, Karlsruhe, Germany), which consists of an array of 47 speakers, or the CUBE at Virginia Tech (US) with 147. Further similar venues are discussed in the conclusion (section 4) of this paper, with regard to both their own characteristic settings and to how to approach the diffusion of acousmatic music in each case.

A grain can be defined as a small particle of sound, typically of a duration between 10 to 50 milliseconds[1] consisting of two main elements: a signal (which can be produced either synthetically or from an already recorded sound) and an envelope, which shapes the signal's amplitude. Gabor [5] called these small particles of sound "a sound quantum" because, when too small, such particles cannot be perceived by our hearing as sound.

In sound synthesis, grains are normally used in big amounts per second, producing rich sound results resembling clouds made of those grains, where the perception of each grain is fully lost, but the effect of grains acting together is not. From the visual point of view, paintings from the pointillist period (mostly those by Paul Signac and Georges Seurat) are a good visual analogy to granular synthesis clouds: from a close perspective, pointillist forms can be seen as exclusively made from tiny little

[1] The duration of grains can vary from case to case, and therefore, durations of 1 to 100 milliseconds can be also considered for this purpose.

coloured dots (similarly to digital pixels); on the other hand, from a farther perspective, forms from those colourful dots are fully revealed. In granular synthesis, the more dense the cloud, the richer the harmonic texture of each sound moment will be. Roads [4,7] classifies sound granulation in different types, mainly the difference between *synchronous* and *asynchronous* granular synthesis. Synchronous granular synthesis works with grains that are all separated by whether the same amount of time, or at least by some type of linear relationship. On the other hand, asynchronous granular synthesis does not present such a strict linear relationship as the synchronous type, and therefore the relationship will typically contain random elements with no linear common elements amongst them as a consequence.

The goal of GS is to translate all of these main features of granular synthesis into the spatial domain *in real time* within a high-density array of loudspeakers. This allows for the most important contribution of GS: its capability to produce flexible times between 1 millisecond to any longer duration among each speaker in the array. The prototypes developed so far utilise mostly *synchronous spatial grains* including clouds, the latter mostly through the usage of diverse and several octophonic settings within the high-density array of loudspeakers. Most prototypes utilise a constant grain duration per loudspeaker, defined by a constant frequency input, which provides both the time that elapses across the entire array of loudspeakers as well as the duration of each spatial grain for each loudspeaker. Hence, and regardless of the actual number of loudspeakers in the array, the principle in which prototypes work is basically the same. Figure 1 below shows single synchronous spatial grains in a 4.0 array. There are however some prototypes which either increase or decrease the grain time between loudspeakers, constantly changing the duration of grains with the entire array of speakers. These are nevertheless synchronous spatial grains, following the definition by Roads already mentioned above [4,7].

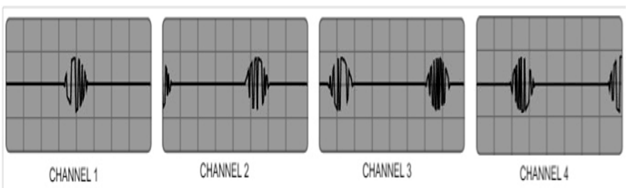

Figure 1. Synchronous spatial grains in a 4.0 array.

Although grains produced by granular synthesis DSP can be indeed diffused in a multichannel environment, the concept of GS proposes to produce the grains in real time, therefore, at the very moment in which the movement between speakers occurs and not before, as the grains are solely the result of the sound diffusion within the loudspeakers array and not synthesised. For GS to happen, a constant signal flow – either a constant, regular signal (such as white noise or a sine wave), or a concrete recorded sound – is required to spatially granulate its output within the multichannel environment. Although spatial grains can be of any size, those specially effective for GS diffusion with a granular aural effect are those which take less than 100ms to travel between any two loudspeakers within any multichannel environment. Moreover, the number of channels of a diffusion array plays a vital role in how the spatial granulation will be not only produced but also programmed, as a large array of loudspeakers within a space can transport the grains across that space much more efficiently and clearly rather than using reduced multi-track systems (such as, for example, a quadrophonic array). Hence, distance can help with the aural recognition of the positioning of the grains, mostly in those cases of spatialisation of very short grains. Figure 2 below shows in this case, that the window for the grain will be 200 milliseconds (for a rotation frequency of 5 Hz), which is the amount of time in which the full four-channel array will be used for the granulation, whilst the spatial grain itself is a quarter of that, in this case 50 milliseconds.

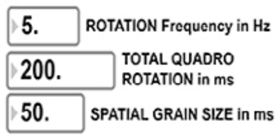

Figure 2. Calculation included in the granular spatialisation software for a quadraphonic granular diffusion. Four spatial grains of 50 ms each will be produced in each of the speakers for a full coverage of the array.

Figure 3. Envelope for grain shaping in a 4.0 surround system, in this case a Gaussian envelope covering only a ¼ of the entire duration of the rotation time window.

Figure 3 above shows how a Gaussian envelope (bell-like shaped envelope, very useful in granular synthesis, see more information later in this section) can be applied to those figures shown in Figure 1 for a 4.0 surround sound system array: the duration of the grain envelope is only a ¼ of the length of the entire window, whilst ¾ of it is silence (which is the time needed for the other three channels to produce the grains). The envelope is thereafter delayed for each speaker by the duration of each grain.

3. GRANULAR SPATIALISATION: CHARACTERISTICS

3.1 Main Characteristics

The diffusion prototypes of the systems programmed at this stage share the following main characteristics:

1. all prototypes have been programmed using the Max 6 software package;
2. The granulation of sound occurs at the specific moment of sound diffusion. Hence, each spatial grain is produced at the precise moment of sound diffusion within a high-density loudspeaker array, not before. Spatial

grains are therefore in charge of the sound's diffusion within the speaker array: the granulation and its aural effects happen in real time at the very moment of spatialisation;

3. One grain per loudspeaker diffusion: the fundamental concept behind GS is the usage of ideally one spatial grain per loudspeaker (which could be even less than one grain, in the special cases of overlapping grains between speakers, explained later in this paper), by exploiting the physical characteristics of grains for each of the speakers within the array, instead of simulating virtual locations as it is the case, for example, with the usage of ambisonics. Hence, each spatial grain has its short specific temporary location in the multichannel system at any given time. This offers a different direction and conception compared to existent development and research in the area, as for example, Scott Wilson's Spatial Swarm Granulation [8], which presents an implementation for dynamic two or three dimensional spatial distribution of granulated sound (therefore, granular synthesis produced in advance of its spatialisation) over an arbitrary loudspeaker system;

4. Granulation effect created within the diffusion: apart from diffusing sound within an *n* number of loudspeakers, GS produces aurally a granulating effect through the diffusion, which varies in its intensity depending mainly either on the type of grain shape or on the overlapping grains between two contiguous loudspeakers or on both characteristics applied together. The diffusion can also comprise clouds of grains, depending on the density of the granulation applied;

5. Different types of spatial grains: although the aim GS is to use any type of grains for the granular spatialisation of sound, so far only synchronous spatial grains were used, mostly due to either their regularity or their linear relationship, which provides for a clear tracing for the signals diffused. Diffusion with exclusively asynchronous spatial grains is also envisaged to be used for either creating dense clouds of spatial grains or for the usage of random duration for the grains. However, at present, both cases need still proper programming and experimentation;

6. Grain-time control between each speaker in the array: the system is conceived to work by controlling the duration (and therefore, the length) of the resulting spatial grains. Although the spatial grain-time could be of any length, for the special case of spatialisation *through* grains, an ideal duration between two speakers would be between ca. 10 ms and 100 ms, although the system uses also both shorter and longer times, the latter in the case of overlapped grains among contiguous loudspeakers. The duration of the grains is defined in this system – as shown in Fig. 2 above – by two parameters:

(a) a rotation frequency (programmed in Hertz), which establishes the time for the spatial grains to cover the entire array of speakers, as defined for each specific case;

(b) the actual number of speakers within the array.

Hence, by a fixed rotation frequency, the higher the number of loudspeakers included in the array, the shorter spatial grains will become. Also, the shorter the grain, the more the typical s*patial granulation* effect can be perceived, including its typical granulated noise;

7. As sounds are constantly and only diffused via spatial grains, the localisation of diverse *spectromorphological*[2] aspects of these sounds – according to Smalley's concept about soundshapes [9] – constantly vary their position in the high-density-speaker array, creating a rather rich and varied s*patiomorphology*, which, also according to Smalley, defines the exploring of spatial properties and spatial changes of sound(s) [9,10];

8. Grain envelopes: spatial grains work in this development with different typical smooth table functions for diverse window envelopes shapes, such as the Gaussian and Quassi Gaussian (Tukey) types, but – depending on the type of granulation desired – they can include sharper envelopes such as triangular, rectangular, etc.;

9. Directionality of the diffusion: spatialisation occurs so far with either a clockwise or an anti-clockwise movement of sound within the arrays of a selected number of loudspeakers. Both directions can use either synchronous or asynchronous spatial grains. As mentioned above, random diffusion and asynchronous spatial grains have yet to be implemented within this development;

10. Grain duration: spatial grains in the prototypes developed and tested so far have either a fixed duration or can also dynamically increase or decrease their duration whilst travelling across an array of multiple speakers. In the latter case, grains either accelerate or slow down the circulation of the granular diffusion. This is a smooth and gradual usage of synchronous spatial grains, as they still possess a linear relationship among them;

11. Flexible speaker array constellation: GS can be applied to any set of multi-speaker diffusion system, from 4.0 to an *n* number of loudspeakers in either multiple arrays of, for example octophonic clusters, or using the entire array at disposal. This allows the system to diffuse sound with a typical spatially designed grain characteristic in particular environments. However, the best arrays are those with around 100 or more loudspeakers, as it is explained later in this article;

12. Overlapping of spatial grains (soothing the granulating effect): although GS has been conceived to produce a granulating effect in the overall outcome, thus, allowing for grains diffusing each particular sound per loudspeaker to be heard as such as well as the sound they transport or diffuse, the concept can also be used to diffuse sound more smoothly by overlapping grains amongst two, three or four loudspeakers. Therefore, the question of overlapping or not spatial grains is relevant with regard to how the effect of granulation should be perceived in the spatialiasation. In its pure, original conception, only one full grain per loudspeaker should be heard, meaning that

[2] Spectromorphology is the perceived sonic footprint of a sound spectrum as it manifests in time [9].

all of the other channels are muted. Through the start and end of the grain envelopes, the effect hereby adds a desired granulation noise the higher the rotation frequency is increased. However, overlapping grains between contiguous loudspeakers soothes the process. In order for spatial grains to be still perceived as such, overlapping neither should be massively long nor should it happen across multiple speakers. Hence, only 2x, 3x and up to a maximum of 5x overlapping should be used hereby, with rather short grain durations, in order for the granulation effect to be still perceived, albeit soother than without overlapping.

3.2 Composing & programming GS for the diffusion of sound at the CUBE's high density speaker's array at Virginia Tech (US).

This section briefly describes the experience of composing and diffusing in concert the acousmatic piece *Spatial Grains, Soundscape No 1* inside the Cube at Virginia Tech in August 2015 during a short residency, mixing together sounds from several and very different parts of the world. The spatial diffusion of this imaginary soundscape takes place exclusively through the usage of GS.

The Cube space has an array of 138 loudspeakers, (a figure that includes 4x subwoofers), all of which are distributed between three floors: ground (64x + 10x), two Catwalks and the grid layer (the latter three, with 20x speakers each, whereas the first Catwalk includes the 4x subwoofers)[3]. Although it is not a huge hall, and moreover and in spite of its name, it is of slightly rectangular form, distances in the plane (for plane waves) are relatively short. However, after experimenting on a daily basis with the system during the residency, distance could be indeed perceived in its 3D constellation and therefore, sounds from either the grid, or from each of the Catwalks were very much identifiable with regard to their location in the space and most surprisingly, even from the four subwoofers in the first Catwalk. In order not to use the entire array of the Cube for each sound – which would have been against both the spectromorphological and spatiomorphological characteristics of those sounds in most of the cases – reduced and located sub-arrays of loudspeakers within the Cube were programmed for the composition. The majority of these sub-arrays are octophonic, with some exceptions, such as a 10x speaker sub-array (stage speakers) in the ground floor and further two sub-arrays with twenty-four loudspeakers each in the first and second floor. This conception allowed for a much more creative and sensible manner of propagating sound, and at the same time, for some sounds to be located at only a restricted area of the entire speaker array, due to, for example, their spectromorphological and spatiomorphological characteristics. The main idea of the piece was to use many different types of sounds, each of which required a player similar to the ones described in figures 5 and 6 below. Each player contains the definition of the type of window for the grain envelope, the direction of the spatial granulation, the rotation frequency – that is, the duration for the grains to complete the n-number of speakers of the sub-array cycle – and the grain size, determined by the division between the rotation frequency and the number of channels within the sub-array.

The players for the spatial granulated diffusion of the piece were included in two main and separated patchers programmed in Max, both of which considered different aspects and possibilities of usage of the Cube's loudspeaker array system. The first Max patcher is based on a mixture of several players, each of which plays only mono files with an output of 8x channels, with the exception of one single player designed for quadrophonic output (for 4x subwoofers in the first floor, front, rear right and left sides), one extra player for the 10x stage loudspeakers placed in a surround disposition on the floor (JBL LSR6328P speakers, different to the other 124), and 2x players using 24x channels each. This first Max patcher did not repeat any loudspeaker in any of the 8x, 10x or 24x settings, as it is shown in figure 4.

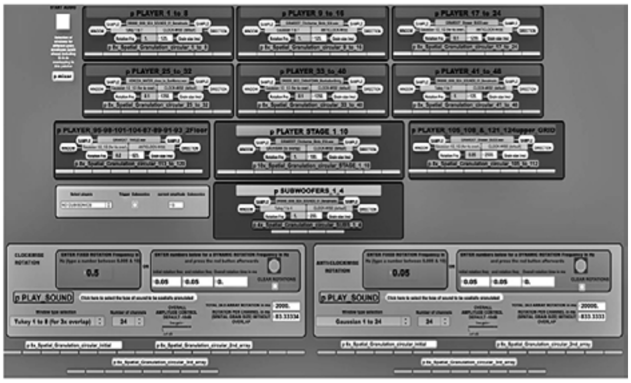

Figure 4. First Max patcher with 4x, 8x, 10x and 24x sub-arrays of speakers for the acousmatic composition *Spatial Grains, Soundscape No 1*.

Figure 5 shows one of the two 24x speakers arrays situated between the upper two Catwalks, with a rather elliptic distribution of speakers, ideal to spatialise sounds of sources such as birds, insects, etc.

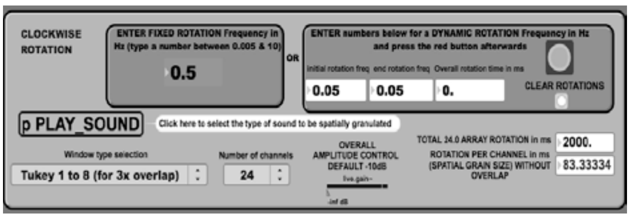

Figure 5. One of the two players of 24x speakers in the first Max patcher. This player can either have a fixed rotation frequency or it can vary constantly the speed of spatialisation, therefore continuously changing the grain size.

Figure 6 below shows *Player_Stage 1 to 10* (the 10x stage speakers on the ground floor), in which the player diffuses the spatial grains within a 10.0 surround sub-array. In spite of this example, most of the players in the first Max patcher have an octophonic surround disposition within each of the floors. The directionality of the granu-

[3] There is actually a total of 147 speakers in the Cube space, thus 9 speakers more than herewith described, but these (9x Holosonic AS-24) were not available during the residency.

lated diffusion can be either clockwise or anti-clockwise for any of the players.

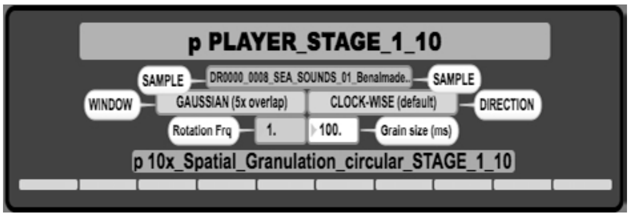

Figure 6. Stage 1-10 player.

The second of the Max patchers features a continuously linear diffusion on all 124 speakers within the entire Cube excluding the 10 stage speakers. Both Max patchers allowed for the usage of the full available array of 138 speakers inside the Cube space.

There was a substantial difference in the disposition of outputs for each of the two Max patchers, with the clear intention of creating different virtual spaces within this large array of speakers, in order for different aspects and elements of the sounds included in the soundscape to become clearly identifiable within the space. With the exception of both 24x spatial grain players as described in figure 5, all of the other players have a surround disposition in each of the three floors with a unique selection of loudspeakers for each player with regard to their exact position within the room. As an example, figure 7 below shows the octophonic surround disposition within the 64 speakers of the ground floor of the Cube of one of the octophonic players. The 8x surround sub-array follows the pattern for which each speaker within the sub-array is equally separated by every seven contiguous speakers, in order to create an individual location for each of the 8x sub-arrays around the audience. Hence, there are 6 octophonic players in the first Max patcher, all of which have a different configuration with regard to their actual speaker numbering (within the Cube, they are numbered from one to 64), whereas no speaker was repeated for any of the players/sub-arrays. The lack of speaker repetition, plus the different settings of spatial grain duration and diffusion direction (clockwise or anti-clockwise), created a fine thread of layers of different sounds and their diverse movements within that particular floor.

The disposition shown in figure 7 below is for *player 33 to 40* only. As mentioned above, all of the other 8.0 players in the ground floor utilise similar combinations and separations for the loudspeakers, without repeating any of the speakers in any of the different combinations. The main reason for such an octophonic surround disposition of speakers in this Max patcher is to avoid that a sound is diffused using only eight contiguous speakers, which would confine that sound to a strict particular and small area of the entire array within each of the floors within the Cube.

With regard to this composition and its performance during the presentation of this paper at ICMC 2016, the main challenge relied on adapting the spatialisation of the composition for a much smaller set of only 16.2 for demonstrations (the submission was originally made to use the 192 speakers of the *Game of Life* system, but the system was not made available for the ICMC in the end).

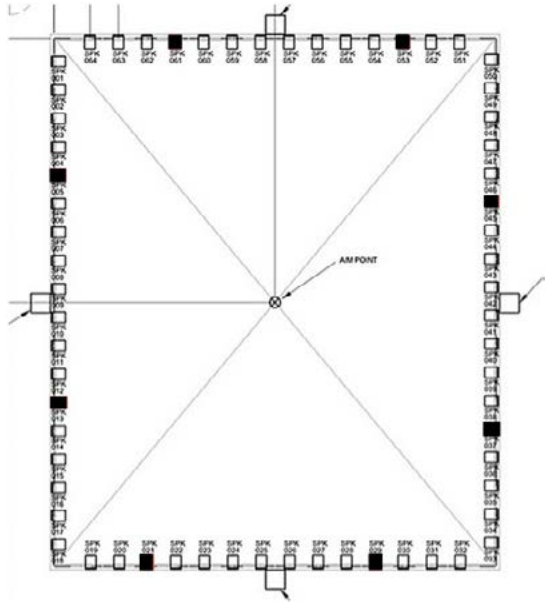

Figure 7. Layout of the ground floor of the Cube space with 64 loudspeakers. The speakers filled in black are those used in one of the octophonic players. The four bigger squares represent the 4x subwoofers.

4. CONCLUSION

The experience in the Cube has shown, that GS can be applied in full for an independent and new manner of diffusing electroacoustic music in high-density arrays of loudspeakers, in spite of the fact, that it can still be effectively used in small diffusion systems such as 4.0, 5.0 or 8.0.

One of the main conclusions after using for the first time GS in a proper high-density array of speakers, is that the Max patchers programmed for the performance of the composition can only be used in the Cube, whilst their translation to any other space and system, such as, for example, the *Game of Life* system in Utrecht (192 speakers), or the BEAST in Birmingham, UK (which counts ca. 100 different loudspeakers) or the *Klangdom* at ZKM, Germany (47 speakers) will fully change the manner in which the piece is performed, and therefore, presents a major challenge with regard to the treatment of spectromorphological and spatiomorphological characteristics of the sounds and their spatialisation within a completely different arranged and conceived environment. From the above, only the Cube's and the *Game of Life's* systems share a basic common WFS[4] conception though.

One of the most relevant differences amongst those spaces mentioned above is related to how timbre works. The Cube is clearly conceived for WFS (in 3D), supported by its equal type of 124 loudspeakers in all three Catwalks

[4] WFS (Wave Field Synthesis) is a spatial sound field reproduction technique performed via a large number of loudspeakers, in order to create a virtual auditory scene over an also large listening area. The concept was initially formulated by Berkhout [11, 12] at the Delft University of Technology at the end of the 80s, based also on previous research, homophony

and the grid. On the other hand, the BEAST in Birmingham, UK, is assembled with several different types of loudspeakers instead of one or two types, featuring as a consequence diverse frequency responses (the most notable, the arrays of tweeters on the roof, to enhance just high frequency sections of the spectrum), and therefore, diffusion of the same piece with this system would present fundamental changes in timbre compared to the Cube's configuration. This implies that, with regard to spatialising the same acousmatic piece in differently designed spaces, several issues must be considered. On the one hand, the programming of the diffusion of pieces must be radically adapted, in order to suit at its best the characteristics of the system used (e.g. the type and number of loudspeakers, their disposition within the space and the variety of applying – or not – different types of loudspeakers inserted within a complex multi-channel array). On the other hand, due to the fact that dramaturgical and timbrical aspects of the pieces can drastically vary from place to place and system to system, the tactics involved in how to use a particular multi-speaker system must be constantly reconsidered, in order to obtain a sound diffusion that suits both the composition and the space with regard to the composition's spatiomorphological contents and potential. This is due to the fact that he usage of different types of loudspeakers in each space may or may not suit the quality of the diffused sounds (including their timbre) and therefore, may or may not serve the intended dramaturgical effect of the composition. With special regard to timbre, GS performed rather equally at every section of the Cube and it proved excellent for carrying sound through the three catwalks and grid with a rather equal timbrical characteristics (most noticeable in those speakers with a WFS configuration). Empirical tests to analyse GS's timbrical response were not possible during the short residency at the Cube, but are planned for the future, in order to fine-tune the system.

From all the above, it is clear that GS is suitable to be adapted for any of those systems already mentioned in this paper as well as other, similar ones, with particular results expected from each different system and space.

As mentioned in the abstract, the number of speakers in the last 20 years seems to only increase: at the Cube in Virginia Tech there are 148 speakers, at IRCAM there are 246, the *Game of Life* has 192. Thus the pertinent question: is the sky the limit for the number of speakers in the future? The ICMC 2016 is asking this question as a conference theme, which seems more than appropriate hereby. With regard to GS, it is clear that it is a new option, which can deal with any number of loudspeakers in order to spatialise sound within any space. In spite of its first full prototypes (as those used at Virginia Tech) were programmed with Max, there is nothing against the inclusion of any other software to do this, such as, for example, Supercollider. In fact, this is my plan for the future of this research. However, Max has proven so far to be a reliable tool for this first and absolved part of this on-going practice-lead research. Thus, depending on the architecture, acoustics, and electroacoustic system design (i.e. the number of loudspeakers in the high-density array) where GS can be applied, new concepts and developments in sound spatialisation and the diffusion of acousmatic music can be explored and expanded.

5. REFERENCES

[1] J. Garavaglia, "Raising Awareness About Complete Automation of Live-Electronics: a Historical Perspective," *Auditory Display*, *6th International Symposium CMMR/ICAD 2009, Copenhagen, Denmark, May 2009. Revised papers - LNCS 5054*. Springer Verlag. Berlin, Heilderberg, pp. 438–465, 2010.

[2] J. Garavaglia, "Full automation of real-time processes in interactive compositions: two related examples," *Proceedings of the Sound and Music Computing Conference 2013* edited by Roberto Bresin, KTH Royal Institute of Technology, Stockholm, 2013, pp. 164–171, 2013

[3] B. Truax, "Real-Time Granular Synthesis with a Digital Signal Processor," *Computer Music Journal* 12(2), pp. 14–26, 1988.

[4] C. Roads, et al. (ed), *The Computer Music Tutorial*. Cambridge: MIT Press, p. 175, 1996.

[5] D. Gabor, Dennis. "Acoustical Quanta and the Theory of Hearing." *Nature 159* (no. 4044), pp. 591-94, 1947.

[6] I. Xenakis, *Formalized Music*. Bloomington: Indiana University Press, 1971.

[7] Roads, C. Microsound. Cambridge, MA: The MIT Press, p. 88, 2001.

[8] Wilson, S. 2008. "Spatial swarm granulation". ICMC 2008 Proceedings, hosted by Michigan Publishing, a division of the University of Michigan Library, http://quod.lib.umich.edu/i/icmc/bbp2372.2008.127/1.

[9] Smalley, D. 1997. "Spectromorphology: explaining soundshapes". Organised Sound 2(2), Cambridge University Press (UK), pp. 107–26.

[10] Smalley, D. 2007. "Space-form and the acousmatic image". Organised Sound 12(1), Cambridge University Press (UK), pp. 35–58. doi: 10.1017/S1355771807001665

[11] Berkhout. A. J., "A holographic approach to acoustic control". J*ournal of the Audio Engineering Society*, 36, pp. 977–995, 1988.

[12] Berkhout. A. J., D. de Vries, and P. Vogel. "Acoustic control by wave field synthesis". J*ournal of the Acoustic Society of America*, 93(5), pp. 2764–2778, 1993.