

A FRAMEWORK FOR THE CHOREOGRAPHY OF SOUND

Gerhard Eckel, Martin Rumori, David Pirrò, Ramón González-Arroyo

University of Music and Performing Arts
Institute of Electronic Music and Acoustics
Graz, Austria

{eckel, rumori, pirro, arroyo}@iem.at

ABSTRACT

A framework developed in the context of the artistic research project *The Choreography of Sound* is presented. The project aims at furthering the practice of electroacoustic composition, especially with respect to the spatial in this genre. The general approach of the project is introduced as a basis for presenting the design of the framework. Its main components are described and the motivations for their design are explained. The framework comprises various tools to formulate spatial aspects of compositional ideas in heterogeneous ways (geometrically, as dynamic systems, and as objects composed from shared properties). Existing spatialisation techniques can be integrated in the framework and exposed to compositional control, such as to allow for unorthodox combinations of them as well as articulating them with other levels of composition. Visualisation and auralisation components allow for sensible work off-site in the otherwise very site-specific project. The *CoS* software framework is conceived as extensions to the SuperCollider audio programming environment.

1. INTRODUCTION

The Choreography of Sound (in the following short *CoS*) is an artistic research project, which aims at furthering the practice of electroacoustic music composition, especially with respect to the treatment of space in computer music. The inquiry into new possibilities of composing spatial sonic entities¹ is carried out through artistic practice, i. e. the composition of case studies, which explore alternatives to thinking sonic space through existing spatialisation approaches. As opposed to conceiving spatialisation as solving a problem of representation, *CoS* understands composing sonic space as a sculptural problem and sound as a plastic object [5] (or even a body) to be composed and choreographed.

Spatialisation is often understood as providing a window into a virtual sound space, which is represented by rendering processes trying to emulate the effects of physical sound propagation. In [1] we can read for instance:

¹The generality of this formulation is motivated by the openness of the research in this project. We cannot know yet which kinds of sonic entities will emerge from our practice-led inquiry. We just know that we do not want to deal only with the notions of sound source objects typical for traditional approaches in sound spatialisation.

“Spatial audio technologies attempt to find a technical solution to create a spatial sound image.” The notion of the image suggests that there are objects which are depicted or rendered audible through spatialisation. An alternative way of thinking about the spatial in music is to compose sound such that spatial sonic entities emerge in the experience of the listener. This does not only involve what typically is understood by rendering, which mainly addresses the sensory aspect of auditory experience. For spatial sonic entities such as *plastic sound objects* or *space-filling textures* (to name only two of those imaginary entities we are interested in) to be experienced, also phenomena on the level of Gestalt perception and cognition (cf. “auditory spatial schemata”, [6]) need to be addressed, i. e. composed.

In *CoS*, we understand composition as extending onto these levels of sonic organisation, i. e. composition also takes place on the level of the spatialisation technique. Therefore it is not enough to include traditional spatialisation parameters in a compositional framework, such as described in [12] as “spatial sound synthesis”. The spatialisation algorithm itself has to be made subject of composition, opened up in order to be linked to any other level of the compositional work. In *CoS*, we are composing possibilities for spatial sound entities to be constituted in and through the enactive perception (cf. [8]) of our listeners. The classical distinction between sound scene description and rendering is not productive in this context, as the entities populating the sound scene cannot be explicitly represented. They are illusions and allusions (cf. [2]) in the body and mind of the listener, which are provoked by compositional decisions on all levels, and not only on the level typically understood as spatialisation.

In our work we are not concerned with rendering sound images for various concert halls using different reproduction systems but with inscribing sound into a specific hall and setup in such a way, that it allows for taking up many different perspectives that expose musical significance. Rather than rendering sound for a *sweet area*, *CoS* attempts to create music which can be perceived from multiple angles or listening positions, each of them provoking a different experience, and all of them forming an integral part of the composition.

As a consequence of this approach, the proscenium type of music presentation may turn out as inadequate: the audience may want to change the listening perspective



Figure 1. Loudspeaker setup in *Ligeti* hall.

during the piece. *CoS* is therefore also exploring different presentation formats mixing concert and installation practices, but these will not be discussed in the context of this paper. This paper rather focuses on introducing the general approach of the project, which forms the necessary background for describing the tools used in the research and for reporting about the challenges of the design and the development of the required software. The general approach of the project is informed by a large number of heterogeneous concepts about the spatial in music and the use of technology for its conception and production, all of which cannot be developed in depth in the context of this paper. Nevertheless, we found it important to introduce at least the most relevant ones briefly.

2. APPROACH

One of the main research activities in the *CoS* project, the composition of case studies, may be compared to the laboratory practice of molecular biologists, as described by Hans-Jörg Rheinberger in his theory of experimental systems and epistemic things [9]. Rheinberger’s experimental systems are “machines to produce future”. They generate the unknown and allow for the unexpected to emerge. Research questions materialise in the experimental system, they are revealed by research practice. They do not necessarily exist at the beginning of the process. The role of Rheinberger’s *epistemic things*, which paradoxically represent what we do not know yet, can be compared to the role of the case study in artistic research.

While composing a case study focusing on a particular

set of compositional problems, the need for certain kinds of *formulations* emerges, informing the parallel process of tool design and development. Furthermore, the metamorphosis of epistemic things, once they are understood, into *technical objects*, which, in turn, become new building blocks in the experimental system, can also be observed with case studies – when the knowledge and understanding they embody forms the basis for new works to be created. Grasping this kind of knowledge is the goal of *CoS*.

Another main research activity of the *CoS* project lies in the construction and maintenance of the experimental system itself. Shortly after the beginning of the project it has been decided to closely attach it to a particular concert hall and its infrastructure, the *György Ligeti* hall of the *MUMUTH* in Graz/Austria (cf. section 3). This decision has been motivated by the need to immerse our practice into one situation and its possibilities, being convinced that only such a concentration may lead to the depth of experience and understanding required to advance our research. We accept the fact that our case studies will probably only be strictly valid in this very context, valuing specificity over generality – an approach frequently taken in artistic research.

The *Ligeti* hall as a central part of the *CoS* experimental system appears, at first sight, as a classical technical object in Rheinberger’s sense. Also the existing spatialisation algorithms *CoS* uses, aiming at combining them in new ways (cf. section 4.4), are technical objects – objects embodying existing knowledge. However, the special infrastructure of *Ligeti* hall, the large number of available loudspeakers along with an advanced speaker positioning system (described in more detail in section 3) also invites to treat approaches to sound projection as epistemic things.

The representation or display components (cf. section 4.5) particularly developed for the *CoS* experimental system have to be seen mainly as technical objects as they are based on existing know-how about three-dimensional visualisation and binaural auralisation, tailored towards the particular requirements of the project. But there are also aspects of these components that are typical for epistemic things, representing what we do not know yet. The display tools allow for certain ways of thinking, imagining and manipulating the main epistemic things of our research – spatial sonic entities. How the display tools will eventually condition the research process cannot be foreseen.

Clearly categorised as epistemic in nature can be the formulation tools developed by *CoS*, which are understood as scaffolds for modelling compositional objects. Among these we find class libraries for representing objects geometrically (section 4.3.1), typically loudspeakers and reflecting surfaces, but also compositional objects with geometric properties, such as sound sources or clouds. Besides relating objects geometrically, there is a need in *CoS* for defining the temporal behaviour of objects organised in dynamical systems (section 4.3.2). For other, more general relations to be modelled, *CoS* pro-

vides an advanced mechanism for dynamically defining and sharing object properties (section 4.3.3).

3. LIGETI HALL

As introduced in section 2, the target place for most of the *CoS* research is the *Ligeti* concert hall of the *MUMUTH*² building in Graz.

The *Ligeti* hall is equipped with a large multichannel loudspeaker setup. 33 speakers³ are mounted on independent, custom made motorised suspensions, which allow for adjusting the height level, the azimuth angle and the tilt of the speakers. Their 2D position with respect to the ground floor area, however, is fixed. Four speakers are mounted in the corners of the room, while the other 29, if put to the correct height levels and angles, can form a hemisphere suitable for optimised Ambisonics soundfield reproduction.

Furthermore, the *Ligeti* hall is furnished with a 64-channel artificial room acoustics system⁴ mounted on the hall's ceiling and on upper parts of the walls, completed by eight subwoofer channels. All channels can be also individually fed, bypassing the acoustics system, which is the primary use case in the context of *CoS*.

Additionally, there are several speakers of different types⁵ mountable on stands at arbitrary positions in the hall, plus two big subwoofers. For some research studies, also the 20-channel IEM Ikosaeder loudspeaker is used [13]. In *CoS*, these speakers are referred to as "flying speakers".

All speaker channels are connected to a versatile mixing and routing matrix incorporating a Lawo *mc*² 66 mixing desk. Two computers, one Macintosh and one Linux PC, are equipped with two MADI interfaces each. Thus each of them can control 128 channels, either individually, in parallel, or in a chain configuration. Typical *CoS* case studies may deal with 30–120 independent speaker channels in total.

The ability of flexible speaker positioning very quickly shifts the roles and understandings of designing and using certain loudspeaker setups in general. Rather than achieving "the ideal" setup mainly based on experience and estimation, the sensual experience and evaluation gains much more importance. It also implicitly devaluates every single setup as it becomes just one possibility out of many easily achievable ones, although it is still the composer's aesthetic and conceptual understanding that arranges the speakers, not the motorised control. As mentioned above, this quality of the *Ligeti* hall partially turns it into an epistemic thing in Rheinberger's sense. It is not merely approached as a fixed technology for sound projection but rather invites and even suggests to regard the methodology of designing loudspeaker setups as an open

²Haus für Musik und Musiktheater, <http://www.kug.ac.at/en/studies-further-education/studies/infrastructure/the-mumuth.html>, last retrieved 2012-02-10

³Speaker type: Kling+Freitag CA 1001-SP

⁴Meyer Sound Constellation Acoustics System with MM 4 speakers

⁵Kling+Freitag CA 1001-SP, Meyer Sound UPJunior, ...

question of artistic research. One result of this are experiments with random speaker positions and orientations in *CoS* as in the case study described in section 4.3.2.

4. FRAMEWORK

4.1. Role of software development in *CoS*

The software framework being developed in *CoS* is conceived as extensions to the SuperCollider audio programming environment. Main reasons for choosing SuperCollider were the program's openness towards different programming paradigms, its performance and multichannel capabilities, its free availability for different platforms and, not of less importance, the fact that everybody in the team already was highly familiar with this software.

During the research process in *CoS*, it became clear quite quickly that the software we want to develop cannot be understood in a one-way functional relationship as a pure "facilitator". This understanding of software as a tool, which serves as a means for realising *a priori* conceived compositional ideas, is quite widespread in the community, even if the influence of these "tools" on artistic processes is widely acknowledged in the same context (as an example, see [1]). As an implication, also software development processes (even those taking place in an artistic context) are rarely considered an intrinsic part of artistic or research activities but rather a separate engineering task of implementing assumed pre-existing, though often implicit, service specifications. This is illustrated by the ever recurring figure of "developing tools for composers in order to make technology accessible" (as e. g. in aforementioned [1]).

In *CoS*, we assume two things which contradict this understanding. Firstly, at least in an artistic context, there is almost never a pre-existing catalogue of "service specifications" which could be implemented e. g. by software. Rather, those "specifications" arise as side effects of an ongoing negotiation between the artists and their means. Therefore, the software development process in *CoS* is regarded as an intrinsic part of and contributing to the artistic research taking place.

Secondly, the artist is not merely the "user", but in a certain sense the "expert" who investigates technology as a facet of his very subject. This includes approaching and incorporating prior knowledge, existing solutions or standard tools, as done by scientists in their respective areas. This understanding is supported (but not coerced) by the fact that in media art, including computer music, the "artist" and the "engineer" often act in a personal union.

Picking up on Rheinberger, this process might be understood in terms of the abovementioned dichotomy between technical objects and epistemic things. In *CoS*, most existing technologies that we adopt are found as technical objects, such as spatialisation algorithms, programming paradigms, software libraries or entire programs. Subsequently, they may become parts of research and creation processes. This means, they may be reshaped or used in different fashions than the originally conceived

ones, and thus transformed into epistemic things. As described in section 2, these research entities represent what we do not know yet but they might generate knowledge by ongoingly “handling” or “dealing with” them. This process might turn back epistemic things into technical objects, allowing to work with them on a different level of understanding.

Applied to *CoS*, this approach has several impacts on the software framework being developed. Not only is it conceived as a library of loose modular building blocks rather than a monolithic application. More importantly, it is not meant to be a tool for certain compositional tasks. In particular, we did not design a framework for describing spatial auditory scenes and rendering them, although it might very well be used as such. Instead, the different parts of the *CoS* software framework should be understood as residues of a thorough examination of approaches to compositional spaces and choreographed sounds.

The explicit understanding of software development in *CoS* as described above does not imply that “classical” software design goals are neglected, such as universality, portability, orthogonality, or performance. They are valid indeed as they are for other software projects. Even more, they also represent technical objects, which play a certain role in the research of *CoS*. From this follows that the *CoS* software framework aims at being universal, with its very much site-specific parts separable as easily as possible from the more generic ones. Most of the code is written from the very beginning with the idea in mind to eventually release independent libraries for the SuperCollider environment.

4.2. Main software building blocks

The research process in *CoS* as described in the previous section so far led to three main building blocks which comprise the current software framework:

- formulation tools, i. e. means for describing spaces as object relations (both abstract objects and relations as well as more concrete ones, e. g. geometric)
- a collection of spatialisation approaches, which can be applied according to objects’ relations
- means of taking up different perspectives, e. g. visualisation and virtual binaural rendering.

4.3. Formulation Tools

One way of composing in *CoS* is to formulate spaces in terms of objects and their relations. The meanings of “object” might be manifold: the term may refer to a spatial sonic entity, be it a “sound source” in the traditional sense or of a more complex form, it may mean an abstract compositional thing or an object in terms of programming. However, for the scope of this section we stick to the general term “object”, as its differing meanings in the various contexts still often refer to the same entity. Likewise, the term “space” may denote various things. A space may

be a purely abstract construct, which organises compositional objects such as musical layers, but also a model of a real space where objects are organised geometrically. Depending on the compositional approach, these two might be partially congruent. Starting from the latter, a universal scene graph concept has been adopted from computer graphics, which is described in section 4.3.1.

Apart from describing object relations in terms of dependencies (hierarchies) and observers using the geometric representation, dynamics modelling allows to construct and evaluate networks of forces between objects, which is described in 4.3.2.

Reflecting the universal understanding of the various roles an “object” might play, its software representation is extended towards a system for the runtime definition, mutation and composition of objects from dynamic property sets. This facility is described in section 4.3.3.

4.3.1. Geometric Representation

As a starting point for *CoS* research, a means for representing the geometric relations of objects was sought. In the first place, the idea was to construct a model of the MUMUTH *Ligeti* hall (cf. section 3), its playback capabilities (loudspeakers) and their layout, and compositional entities such as sound sources. This should serve as a generic basis for the application of common spatialisation approaches. For many standard or non-standard spatialisation algorithms, different parameters of those geometric relations need to be computed, such as the distance of a sound source to each of the speakers for DBAP (*Distance Based Amplitude Panning*, [7]).

On the other hand, it might be desirable to express geometric object relations in hierarchic structures. This may include the grouping of certain speaker arrangements to a single entity, as in the case of the 20-channel IEM Ikosaeder speaker (cf. section 3), but also modelling spatial dependencies of sound objects, such as multi-source objects or swarms.

Based on these findings, as a generic representation for geometric relations a scene graph model of three-dimensional hierarchic coordinate systems has been implemented. As opposed to computer graphics, the scene graph approach here is not so much a means of optimisation for rendering but rather of structuring conceptual relations. Furthermore, it provides the basis for a largely uniform interface to different spatialisation techniques.

An initial attempt of including a fourth (e. g. temporal) dimension with the geometric representation was dropped so far. It turned out that the various modes of dealing with time in music composition cannot be easily subsumed under an extended geometric understanding without substantially narrowing the practical access to temporal processes.

The adopted scene graph approach within the *CoS* framework is one example for an existing technical object, adopted from computer graphics, and subsequently turned into an epistemic thing. This object is characterised

by the confrontation of well-known computer graphics technology with a possible way of formulating computer music, combined with an expectation that this might reveal knowledge on spatial composition which has been only implicit before.

4.3.2. Dynamics Modelling

From the experience gained in the first *CoS* case studies, the necessity emerged to extend the static, geometric links between objects by relations that could alter the objects' properties or their state over time, i. e. dynamically.

In our thinking, these alterations are caused by effects or "forces" objects exert on each other. In the dynamics modelling subsystem every object contains a property describing how its effect applies to others and a record of all objects affecting it. Objects that are interconnected by these forces form a *system*. Systems evolve by triggering all the effects and applying the changes to all objects at once, setting them into a new state. Furthermore, each system advances at its own pace, allowing to have parallel systems developing at different rates.

For instance, with these extensions, particle-based physical models can be implemented and simulated. In this case, the properties on which forces act are the objects' position, velocity and acceleration. Their successive states define trajectories in a three-dimensional space.

Thus, objects are linked not only by their geometrical disposition, but also by the influence they have on each other i. e. by the interaction taking place between them. The dynamics modelling subsystem provides means to formulate interaction mechanisms, making these subject to composition. Changing the definitions of their effects means to change how objects act and react, how they behave with respect to the others. Using different forces implies that objects inscribe different paths in space and time making different *dynamics* of their movement appear. Composing the interaction between objects allows to model dynamics and behaviour, eventually linking the spatial and temporal appearance of sound in a choreography.

To illustrate how dynamics modelling is employed in the *CoS* project, we describe one of the case studies using this approach. For this experiment all the motorised loudspeakers of the *Ligeti* hall are placed and oriented irregularly. We use the *CoS* framework to represent them geometrically and group them in a dynamical system. In this system, the loudspeakers act as if they were "turbines" (or fans), attracting objects towards them from one direction and throwing them out in the direction they are oriented to. Further, reflecting walls, floor and ceiling have been added to constrain moving objects to a space corresponding to the actual hall. This approach is motivated by the desire to link the features of the loudspeaker configuration (positions and orientations of the speakers) to the trajectories produced by the dynamics system. An ADBAP object (cf. section 4.4.1), injected into this space, will move according to the force fields of the turbines it encounters on its way. The ADBAP algorithm "sonifies" this motion by

spatialising the sound source along its path. The dynamical system generates movement, a spatial and temporal behaviour, eventually choreographing the appearance of sound in the *Ligeti* hall.

By this example two perspectives on dynamics modelling can be made evident. From one point of view, interactions take effect on the sound source, pushing it into movement. From the other, force fields form and shape the space in which the sound source lives. Here, movement is the result of the source "exploring" valleys dug and mountains raised into its space by the forces, trying to reach its equilibrium. In practice, both points of view are equivalent.

4.3.3. Generic Objects with Hierarchic Properties

Finding a representation of more generic objects is not as straightforward as in the case of geometric objects or those designed for dynamics modelling. As different kinds, uses and functionalities of such objects are expected to arise during the research process, it became clear that a flexible object framework is needed that can be extended and adopted according to the developing needs and applications.

The current attempt of such a framework are objects that are characterised by a set of *properties*. These are basically dictionaries of key/value-pairs, which may contain both primitive (e. g. numbers) or complex objects as well as functions. Property dictionaries can be used locally for a single generic object, but they may be also shared among multiple objects. As property dictionaries involve a notification mechanism to the objects they are referenced by, this can be used to update multiple objects upon changes of shared property sets.

Property dictionaries may in turn contain property dictionaries, such that hierarchically structured property sets can be constructed. On any level of such a hierarchy, a property dictionary may be shared or private. Even subtrees of shared dictionaries can become private again, which is made possible by a sophisticated yet entirely transparent per-object-handling of subtrees.

The *CoS* property system is completed by convenience functionalities such as read-only property sets or optional copy-on-write semantics for an automatic transition of shared dictionaries to private ones.

The open design of the *CoS* property system allows for specifying objects using different design patterns in a fully dynamic way. Template property subtrees may be referenced and further specialised in order to furnish an object with certain functionalities. This is similar to object oriented concepts like multiple inheritance, interface/protocol definitions or object composition whose application happens only at runtime.

In *CoS*, several building blocks of the software framework are interfaced to the objects by such sets of properties. For example, the parameters that determine the way how an object is graphically represented are conceived as templates for property subtrees (cf. section 4.5.1). Cer-

tain spatialisation algorithms that evaluate further data beyond geometric coordinates are fed by property dictionaries (section 4.4), as well as objects' characteristics in the context of dynamics modelling (section 4.3.2).

4.4. Spatialisation Algorithms

In *CoS*, most spatialisation algorithms operate on the level of the geometric representation. Usually implemented as a geometric object itself, a spatialisation object is associated with a set of other geometric objects. For example, in the case of DBAP, the associated objects represent the speakers whose distances are used to calculate the amplitude weights, while the DBAP object itself represents the position of the sound source. Other algorithms may take into account further objects such as listeners, absorbers or alike.

As geometric objects are subclasses of generic *CoS* objects, they also provide hierarchic properties (section 4.3.3). Spatialisation algorithms may evaluate subtrees of these property sets in order to allow for specific functionality that cannot be controlled by geometric coordinates only, such as radiation patterns. Even entirely abstract (non-geometric) objects may be used, e. g. for controlling larger groups of spatialisation processes.

The integration of spatialisation algorithms with the objects they are related to allows for easily replacing or manipulating spatialisation processes. Also, many of those may happen at the same time, using independent or intersecting groups of speakers, sources or abstract compositional objects. The combination of a three-dimensional geometric representation and the versatile hierarchic property system facilitates the efficient evaluation of standard spatialisation algorithms while advanced techniques operating on more complex data may still be conceived within the same framework.

However, this uniform interface to different spatialisation algorithms must not be misunderstood as a design decision towards a separation of an auditory scene description (sources, trajectories etc.) and the spatial rendering, as claimed by many spatialisation software projects or standardisation attempts (for an example, cf. [12]). Although it is possible to use the software framework following this paradigm, and although in some cases of spatial composition this might be meaningful, in *CoS* the spatialisation is regarded as an intrinsic part of musical composition which generally cannot be separated from other kinds of sound synthesis or -processing (cf. section 1).

This is underlined by the fact that spatialisation algorithms in the *CoS* framework do not operate directly on the digital signal processing level but are rather used as parameter generators for a separate signal processing stage. It is left entirely open, how and at which level of abstraction these parameters are applied. For many spatialisation techniques, however, the signal processing actually taking place boils down to amplitude weighting and possibly the application of delays and/or simple filters.

The spatialisation algorithms in the *CoS* software so far are implementations of well established sound dif-

fusion techniques such as DBAP and a variant (section 4.4.1), a simplified WFS approach (section 4.4.2) and prospective interfaces to Ambisonics spatialisation (section 4.4.3).

It is, however, not the goal in *CoS* to have a complete suite of spatialisation algorithms. On one hand, those algorithms can be easily implemented the "soft way" without appearing as an object class of its own, e. g. using properties and notifications or simply SuperCollider functions. This is actually the case for evaluating several non-standard algorithms. On the other hand, extending the library of geometric spatialisation objects is very easy, such that they can be added on demand when they are needed.

4.4.1. DBAP and ADBAP

As the DBAP amplitude panning algorithm makes no *a priori* assumption of the effective loudspeaker setup, this spatialisation method was a sort of natural choice in the *CoS* case studies as it best fits the project's general approach.

The DBAP algorithm implemented in the *CoS* framework is a three dimensional extension of the technique formulated in [7]. Having specified the rolloff and blurring coefficients, the distances of the source to the associated loudspeaker objects are computed and used to determine the relative amplitudes of the projected sound.

Basing on the principle of constant intensity panning, the original method requires that the squares of the amplitudes sum up to unity. The resulting overall intensity is constant, regardless of the position of the source. As a consequence, positions outside the loudspeaker field cannot be clearly rendered: in this region the relative amplitude differences tend to zero with increasing distance while the overall intensity is still kept constant, resulting in a spatially undifferentiated sound output.

However, in the course of our case studies, it turned out necessary to spatialise sources that could also travel out of the loudspeaker field and completely disappear. To achieve this, we modified the DBAP algorithm removing the constant intensity condition: sound spatialisation is achieved defining a distribution of absolute rather than relative amplitudes. This causes sources that move sufficiently far away from the loudspeaker array to fade out.

Furthermore, we discovered that especially the trajectory of moving sounds (as in the example described in section 4.3.2) appears more clearly shaped or "sharper", compared to the unmodified DBAP algorithm.

Lacking a more explanatory name, we call the simplified version of the DBAP algorithm Absolute Distance-Based Amplitude Panning (ADBAP).

4.4.2. WFS

At this point the *CoS* framework offers a simplified implementation of the Wave Field Synthesis (WFS) technique (cf. [3]) for experimental purposes. It supports only the calculation of time delays for plane wave synthesis, as a

function of the plane wave orientation and the speaker positions. It is noteworthy that the implementation of the algorithm based on the underlying geometrical representation takes only four lines of code. This illustrates that existing spatialisation techniques can be added easily and whenever demanded from a compositional point of view. The structure of the framework supports the conception and implementation (i. e. composition) of new techniques and the experimentation with them – especially if they exploit geometrical relationships between loudspeakers and spatial sonic or compositional (i. e. conceptual) entities. In the process of developing *CoS* case studies, this WFS implementation was used to perceptually estimate the effects of irregularly spaced loudspeaker arrays exhibiting spatial alisasing due to their sparsity (cf. [4]).

4.4.3. Ambisonics

So far, spatialisation with Ambisonics does not play a major role in *CoS*. This may be derived from the principal research approach as described above, since Ambisonics is a holophonic soundfield reproduction technique, which is restricted to a sweet spot or area. However, it is still desirable to have Ambisonics available for approaching it in an experimental way using subgroups of closely related speakers, under-determined or non-ideal speaker layouts, or making various, often hidden parameters of Ambisonics encoding and decoding available for composing.

A basic interface to existing Ambisonics implementations such as *AmbiEM*⁶ can be easily realised within *CoS*: geometric parameters like relative orientation of objects (i. e. azimuth and elevation) are available from the geometric representation (cf. section 4.3.1) while there are no assumptions on the actual signal processing.

For further experimentation it is desirable to have access to advanced Ambisonics features, such as mixed-order systems, distance coding or multi-band decoding. Prior research has been undertaken to design such a framework for SuperCollider, called *Girafe* [10], that benefits from both the powerful language and from the multichannel routing capabilities of the server. Promoting a separation of signal processing and control data generation itself, it nicely integrates with the design principles of the *CoS* framework, but its integration is not yet finished.

4.5. Display Tools

4.5.1. Visualisation Subsystem

For taking up different perspectives, a visualisation subsystem with interaction capabilities is being conceived. It consists of a precise three-dimensional model of the *Ligeti* hall including the speaker lifting mechanism and models of the other speakers used in the project. For displaying and navigating the model, the *Blender*⁷ game engine is

used, which is controlled via OSC⁸ from the *CoS* software system in SuperCollider.

The visualisation system of *CoS* attempts to fulfil several functions. Firstly, it serves as a kind of display and debugging tool by visualising the geometric layout of certain constellations defined on the textual side of the *CoS* system. It is also used for creating documentary screen shots or videos.

As the access to *Ligeti* hall is limited, the visualisation system also serves as a means for working on case studies outside the actual space. Due to its common acceptance, there is an established understanding on how a three-dimensional visualisation of a space relates to the pictured reality. In *CoS*, we assume that this relation can be exploited to support the composer's abstract imagination of a certain space, which works on the level of experience and retrospection rather than that of a simulation.

In another way, the visualisation system is also used as a construction tool, e. g. for designing new loudspeaker setups (cf. figure 1). Some of the setups used in *CoS* case studies are designed by composing first-order wall and floor reflections of imagined on-axis speaker radiation at the same level as direct sound from the speakers. The model here serves for visualising and estimating these propagation paths in a simple way. This must not be mixed up with fully fledged auralisation software or room acoustics simulation tools, which would be far beyond the needed degree of detail of *CoS*.

4.5.2. Virtual MUMUTH – Binaural Auralisation

Another display in the *CoS* framework is provided using binaural projection of the measured *Ligeti* hall acoustics. As for the visualisation subsystem, the motivation behind the binaural auralisation is not a simulation allowing for a perfect immersion, instead and again, it is meant to take up a different perspective on the acoustic and auditive processes taking place in the real space. It is therefore also an investigation on the viability and validity of using such simulation tools for the process of composing music.

Virtual MUMUTH is realised using sets of measured binaural room impulse responses (BRIR) in a multichannel convolution matrix [11]. Every set of those impulse responses corresponds to only one loudspeaker setup and only one listening position and orientation. Therefore, no advanced techniques from binaural room simulation systems can be applied, such as rotation compensation using head tracking, and the auralised configuration is mostly static. On the other hand, using BRIRs preserves the actual acoustic properties of the *Ligeti* hall in the sense of communications engineering to a much larger extent than common room simulations. The direct comparison of the binaural auralisation and the real acoustics in informal on-site listening tests showed that a sufficient *suspension of disbelief* can be achieved, such that both are not anymore clearly distinguishable. Thus, the restriction to selected static configurations of speakers and listening positions is

⁶<http://sonenvir.at/downloads/sc3/ambiem/>, last retrieved 2012-06-01

⁷<http://www.blender.org>, last retrieved 2012-02-20

⁸<http://opensoundcontrol.org>, last retrieved 2012-02-20

not regarded as a flaw in *CoS* but it is taken as the constraint of an otherwise very efficient tool. Starting from this observation, our investigations follow questions like “As I have a certain memory of the real acoustics by experience, how can binaural auralisation allow for recalling this memory and for estimating how new material may sound in the real space?”

The binaural room impulse responses used in *CoS* include the full reverb tail of the *Ligeti* hall, therefore they have a length of 65536 or even 131072 samples (1.4 or even 2.8 seconds at a sampling rate of 44100 Hertz). Despite this large size, it is still possible to perform a realtime convolution for full speaker sets of 120 channels on fairly recent multi-core machines using the efficient convolution engine *Jconvolver*⁹.

5. CONCLUSIONS

Although the development of the *CoS* software framework is not completed yet, the experiences made with the production of case studies so far has shown that the decisions the development has been based on (cf. section 4) are productive. Being able to represent objects geometrically has led to new ways of thinking the spatial in electroacoustic music, especially when using a concrete loudspeaker setup in a particular hall. Including the geometrical representation of a (pre-composed) loudspeaker setup and the reflecting surfaces in the hall into the compositional model helps approaching one of our epistemic objects – realising a quasi-plastic sound matter. Making such sound matter subject to choreographic treatment, i.e. compose its movement in space, can be approached successfully by organising sonic objects in dynamical systems, including the loudspeaker setup and other relevant objects in the hall (e.g. reflectors). The binaural auralisation possibility, together with the visualisation of the speaker setup allows for a sensible work outside of the *Ligeti* hall, given that one has experienced the respective case study in the hall at least once before listening to it through the auralisation.

The *CoS* software framework will be made freely available as an extension library to SuperCollider on the *CoS* website.¹⁰ Additionally, the three-dimensional model of the *Ligeti* hall, several sets of binaural room impulse responses and case studies will be available at the same place.

6. ACKNOWLEDGEMENTS

The *CoS* project was conceived by Gerhard Eckel and Ramón González-Arroyo. Together with Martin Rumori and David Pirrò they form the core team of the project at IEM in Graz, maintaining links to researchers at BEK in Bergen, CIRMMT in Montreal, ICST in Zurich and IRCAM in Paris. Thanks to Jonas Hansen for 3-D modelling of the *György Ligeti* hall. *CoS* is funded by the Austrian Science Fund (FWF): PEEK AR41.

⁹<http://kokkinizita.linuxaudio.org/linuxaudio/>

¹⁰<http://cos.kug.ac.at>, last retrieved 2012-05-23

7. REFERENCES

- [1] M. A. J. Baalman, “Spatial composition techniques and sound spatialisation technologies,” in *Organised Sound 15*, 209–218, Cambridge, England, 2010.
- [2] N. Barrett, “Spatio-musical composition strategies,” *Org. Sound*, vol. 7, no. 3, pp. 313–323, 2002.
- [3] A. J. Berkhout, “A holographic approach to acoustic control,” *J. Audio Eng. Soc.*, vol. 36, no. 12, pp. 977–995, 1988.
- [4] E. Corteel, “On the use of irregularly spaced loudspeaker arrays for wave field synthesis, potential impact on spatial aliasing frequency,” in *Proc. of the 9th Int. Conference on Digital Audio Effects (DAFX-05)*, 2006.
- [5] R. González-Arroyo, “Towards a plastic sound object,” in *Raum: Konzepte in den Künsten, Kultur- und Naturwissenschaften*, P. Ernst and A. Strohmaier, Eds. Baden-Baden: Nomos, 2012, forthcoming.
- [6] G. S. Kendall and M. Ardila, “The artistic play of spatial organization: Spatial attributes, scene analysis and auditory spatial schemata,” in *Computer Music Modeling and Retrieval. Sense of Sounds*, R. Kronland-Martinet, S. Ystad, and K. Jensen, Eds. Berlin, Heidelberg: Springer, 2008, pp. 125–138.
- [7] T. Lossius, P. Baltazar, and T. de la Hogue, “DBAP – distance-based amplitude panning,” in *Proceedings of the International Computer Music Conference*, Montreal, Canada, 2009.
- [8] A. Noë, *Action in Perception*. The MIT Press, 2004.
- [9] H.-J. Rheinberger, *Toward a History of Epistemic Things: Synthesizing Proteins in the Test Tube*. Stanford, CA: Stanford University Press, 1998.
- [10] M. Rumori, “Girafe – a versatile ambisonics and binaural system,” in *Proceedings of the Ambisonics Symposium 2009*, Graz, Austria, 2009.
- [11] M. Rumori, F. Hollerweger, and A. Cabrera, “Binaural room impulse responses for composition, documentation, virtual acoustics and audio augmented environments,” in *Proceedings of the 26th VDT International Tonmeister Convention*, Leipzig, Germany, 2010.
- [12] M. Schumacher and J. Bresson, “Spatial sound synthesis in computer-aided composition,” in *Organised Sound 15*, 271–289, Cambridge, England, 2010.
- [13] F. Zotter and A. Sontacchi, *Icosahedral Loudspeaker Array*. IEM Report 39/07, 2007. [Online]. Available: http://iem.kug.ac.at/fileadmin/media/iem/altdaten/projekte/publications/iem_report/report39_07/report39_07.pdf