utruchirp - An Impulse Response Measurement and Auralisation Tool Developed for Artistic Practice

Ludvig Elblaus
elblaus@kth.se
Sound and Music Computing Group
KTH Royal Institute of Technology
Stockholm, Sweden

Gerhard Eckel
eckel@iem.at
Institute of Electronic Music and Acoustics
University of Music and Performing Arts Graz
Graz, Austria

ABSTRACT

This paper presents the *utruchirp* software, a tool for measuring impulse responses and modelling room acoustics in real time through auralisation based on convolution using those responses. utruchirp is the result of concerns and needs emerging from the authors' ongoing artistic practice, exploring room scale acoustic feedback as material for live performance, installations, and fixed media pieces as *utrumque*.

The paper provides the technical and, more importantly, the artistic details of the development of utruchirp and its features, highlighting those that are the direct result of insights from artistic work: Monitoring of all stages of measuring and signal processing, auralisations of the measurements from within the measurement process, and integrated round trip delay estimation. Finally, it points out future directions and features that are to be explored next, with an invitation for collaborative efforts, aiming to bring the sensibilities of musical instruments to our measurement tools.

CCS CONCEPTS

• **Applied computing → Performing arts**; • **Software and its engineering → Designing software**.

KEYWORDS

impulse response measurement, acoustics, convolution, room modelling, feedback, composition, electroacoustic music.

ACM Reference Format:
Ludvig Elblaus and Gerhard Eckel. 2020. utruchirp - An Impulse Response Measurement and Auralisation Tool Developed for Artistic Practice. In *Proceedings of the 15th International Audio Mostly Conference (AM'20), September 15–17, 2020, Graz, Austria.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3411109.3411140

1 INTRODUCTION

In this paper, we propose an approach for how to use measurement and modelling of room acoustics for the explicit purpose of artistic practice and artistic research. More precisely, in the first part of the paper, following the introduction, we briefly describe the swept-sine room measurement technique, room acoustics simulation in general and our chosen method in particular. From this, in the second part of the paper, we introduce our new software tool, *utruchirp*, showing how our practice is inscribed in the structure of the tool itself and what experiences motivated the major design decisions. The tool will be shared as open source at the time of publication of this paper. In the third part of the paper we propose future directions in the artistic use of room acoustics simulations and the further development of utruchirp.

utruchirp is a tool primarily for estimating the acoustic responses of rooms, although any acoustic system can be measured. For that purpose the rooms are exited with loudspeakers and their responses are captured with microphones. Special test signals are used, making it possible to compute so-called *room impulse responses* from the captured microphone signals. The measured room impulse responses portray the acoustics of the room under the specific perspective of the transducers involved in the measurement, i.e. under the constraints of their placement in the room and their characteristics such as radiation and directionality patterns.

The impulse responses can be used to create a linear model of the behaviour of the whole electroacoustic system composed of transducers and room. The simulation of the system behaviour is performed through convolution and is generally referred to as *auralisation* [19]. A sound signal convolved with the measured room impulse response will be suffused with the characteristics of the whole system, i.e. transformed as if it was projected by the loudspeaker into the room and captured by the microphone used during the measurement.

It should be noted that the contribution of this work is not the technical methods used for measurement and simulation, i.e. the swept-sine technique and convolution. They are for the purposes of this paper considered as solved problems. Rather, it is the particular presentation and design that allow these techniques to be integrated successfully into a musical practice, highlighting immersion and a sustained experiential perspective throughout an otherwise technical measurement process. As such, the work can be framed as artistic research, and the results an artistic contribution rather then an engineering one. In addition to this, the artistic contribution has not only been extracted from our practice for the purposes of discussion and sharing, it has been reformulated as open source software that we invite others to engage with, put to use and modify. Therefore, some details of the software might change after this publication, but the fundamental goal of the approach will remain, to have measurement tools that feel like musical instruments.

AM'20, September 15–17, 2020, Graz, Austria

ACM ISBN 978-1-4503-7563-4/20/09…$15.00
https://doi.org/10.1145/3411109.3411140

2 BACKGROUND

Using models based on impulse responses to create simulations of real acoustic environments can be of high value for artistic applications that specifically exploit the characteristics of particular combinations of loudspeakers, microphones, and spaces. This is the case with our practice, which focuses on the site-specificity of acoustic feedback processes. Once a measurement of a setup to be explored artistically exists, large parts of our usually site-specific compositional work can be performed off-site in simulations, thus not requiring physical access to the setup in question.

Beside the practical benefits, the systematic work with simulations allows us to greatly refine our practice of engaging with feedback systems. By modelling complex multi-transducer or even multi-room sites, we can explore the particularity of an acoustic situation in unprecedented depth. In a typically highly iterative empirical process, we provoke and explore the complex acoustic behaviours we are seeking in our work. A prerequisite for learning from such explorations is the stability in the reproduction of the results, which can be guaranteed because the room and transducer related parameters are fixed in the simulation. Furthermore, by comparing different measured sites, we can get a sense of which aspects of the found or induced behaviours are linked to the site and which to the process we composed.

Reproducibility is also an important requirement for our collaborative compositional work, especially when working at distance, which we increasingly do. The possibility to share experiences in simulated spaces is therefore the basis for our co-creative practice. It is not unusual for our projects to be developed and rehearsed in simulated environments for weeks or months before moving into the actual venue for a performance. For more details on this aspect of our work, we refer to our paper *Acoustic Modelling as a Strategy for Composing Site-Specific Music* [6].

The risk of overfitting our work to a particular simulation can be efficiently reduced by testing our approaches in different simulations. Therefore it is of great interest to dispose of a large database of room impulse response measurement, ideally with systematically varied transducers setups at the same site. utruchirp is made to allow us to expand this database with as little effort as possible. Measurements of the very spaces we performed in also enable us to produce high-quality fixed-media releases from multi-channel recordings of our concerts or create work entirely within simulations, opening possibilities not available in physical spaces.

The fundamental measurement technique for the work in this paper is known as *swept-sine measurement* procedure and is a standard system identification method widely used in acoustics research [7, 8]. Although there are alternatives to the swept-sine technique, e.g. those based on Maximum Length Sequence pseudo random signals, it is by far the most robust approach known today. It results in a good signal-to-noise ratio, deals with and separates out nonlinearities, and does not require playback and recording sample clocks to be strictly synchronised.

Because of it simplicity and efficacy, convolution is extensively used in media production for sound reverberation effects, e.g. musical production, film post production and video game systems [16]. Rooms with good or interesting acoustics are measured and the reverberation models produced are applied to sound signals that are to be reverberated through convolution. It is also the simplest way to accurately simulate a particular space, if for instance a video game takes place in a real world location with a particular reverb, or if a musical recording is done in a space and later overdubs need to blend with the takes recorded on site. In our practice however, as stated in the previous section, the simulation of the interaction of particular loudspeaker/microphone setups with acoustic spaces is mainly used to experiment with dynamical feedback systems. Room simulation using convolution is a methodology that offers site specific electronic music, like our feedback works, the same comfortable digital laboratory that the studio of an electroacoustic composers does. Complex configurations can be set up, recalled, and compared in an instance. This allows us to work in systematic and reproducible ways that would be very hard to realise in a physical space using large numbers of full range speakers and microphones, for practical reasons.

The challenges of diffusing music on physical setups in a live environment persist, of course, both for fixed media pieces and pieces that use room scale feedback, as well as all other music performed on or amplified by loudspeakers. Yet, the benefits of preparing and composing in simulated acoustics have proven invaluable to us. Therefore, since this approach is central to our practice, we need to be able to estimate room impulse responses reliably and with as little effort as possible.

2.1 Measurement Tools

There are many commercially available tools for room acoustics measurements. Many of them are able to perform the swept-sine technique to estimate room impulse responses, see for instance *DIRAC* from *Brüel & Kjær*, *IRIS* from *Marshall Day Acoustics*, or *Smaart* by *Rational Acoustics*. Most of the commercial options offer a wide variety of measurements and analysis methods and many of them produce standardised metrics for the evaluation of acoustic environments. In other words, they are not focused on measuring and simulating acoustics for musical production purposes, and being closed source it is impossible to modify them, making it difficult to adapt them to specific artistic needs and applications.

There are two exceptions to this we want to mention here: the impulse response measurement tools *ScanIR* for *MATLAB* by Boren et al. [3, 18] and *ALIKI* by Adriaensen [1], which are both are open source. In our practice, we have so far used a version of ALIKI, compiled for macOS by Martin Rumori in the context of the artistic research project *The Choreography of Sound* (CoS) [5], that is adapted to measure large setups.

The development of utruchirp is based on comprehensive process knowledge and practical experiences with measuring room impulse responses and performing auralisations gained in CoS. One of the results of CoS was the application *StiffNeck* developed by Rumori, which enables structured access to a large amount of loudspeaker/microphone configurations in the CoS room impulse response database [4]. Measured configurations can be combined into auralisation setups using the *Jconvolver* program by Adriaensen [2]. The CoS database was created by Rumori using a tool chain based on ALIKI [9]. We used a subset of the CoS tool chain until the possible improvements of this approach to further align it

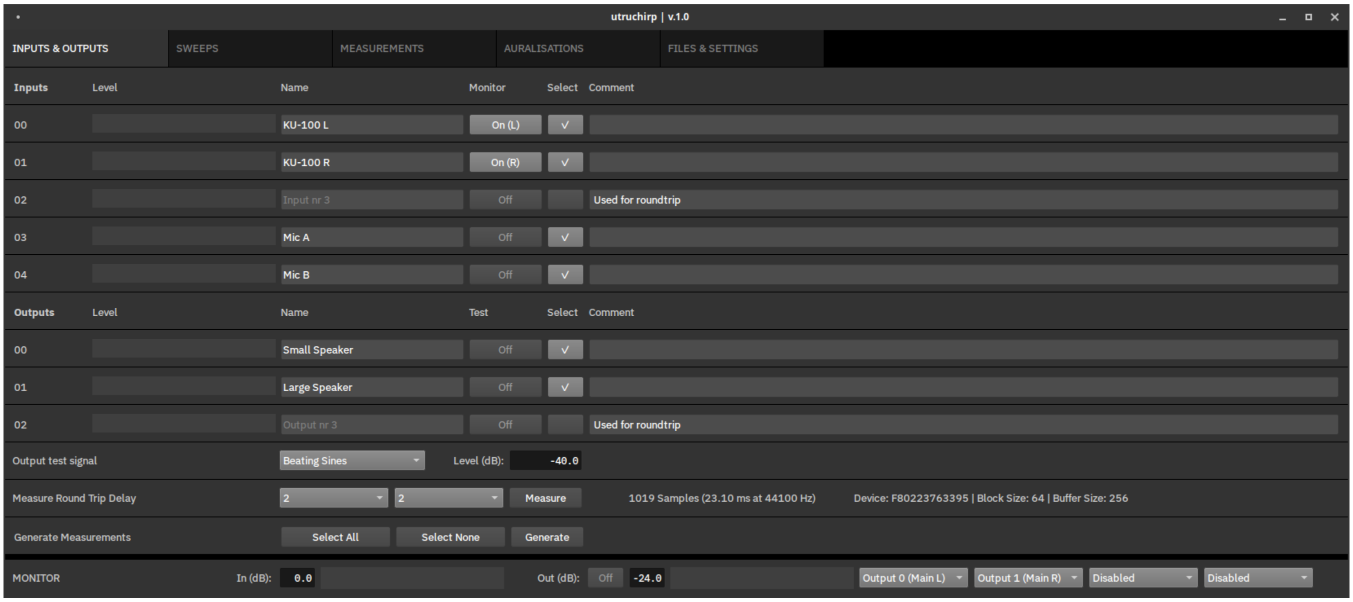

Figure 1: The utruchirp graphical user interface.

with the practicalities of our work where sufficiently understood to design utruchirp.

3 UTRUCHIRP

In this section, we describe how utruchirp is put together, present a selection of the features it provides and describe how a typical measurement and auralisation workflow might look like. We aim to provide enough detail to give a clear sense of the process of setting up a measurement, measuring a space, and creating an auralisation, but several parts have been left out for the sake of brevity. Some features are described in further detail in Section 3.2–3.4, as they show how the particular concerns of our compositional practice have been inscribed in the design of the tool.

utruchirp consists of a set of classes for the SuperCollider programming language, originally by McCartney [14, 15], and more recently described in *The SuperCollider Book* [20]. It does not rely on any extensions to sclang, the SuperCollider interpreted programming language, or plugins to the SuperCollider server. However, it relies on Jconvolver, a separate program that performs the real-time convolutions and provides a Jack [11, 12] interface for routing audio in and out of the convolution process. This might change in future versions, and is further discussed in Section 4.3. But for now, to run utruchirp, the requirements are SuperCollider, Jconvolver, and Jack.

SuperCollider was chosen as a platform to develop and run utruchirp because it is open source, with a GNU GPL license, and it is cross-platform compatible with practically no configuration. Furthermore, being written in sclang, one can easily change or extend utruchirp without having to setup build environments, keeping the threshold of entry for external contributors low. This attitude is also reflected in the code base, where it has been a design goal to keep everything modular, with as small indivisible components as possible, making it easy to pick the system apart, to reassemble it in other configurations, and extend it.

When measuring room impulse responses, quality assurance is very important, as many things can go wrong during a measurement. Often measurements can only be performed at particular times and may possibly not be repeated for various reasons, such as reduced availability of the room and/or audio setup in question or disturbances through noise. Therefore utruchirp supports the user in assuring that each step of the measurement is successful. This concerns making sure that the signal routing is correct, the playback and recording levels are adequate, the obtained signal-to-noise ratio is sufficient, and the measurement is free of disturbing sounds and resulting computational artefacts.

The major features of utruchirp we will foreground in this section are: 1) Reliable monitoring of input and output signals in the measurement session, 2) visual and auditory verification of intermediate measurements results during the measurement process, 3) automatic handling of auralisation configurations, 4) visual and auditory comparison of different measurements and transducer configurations, and 5) integration of automatic I/O round trip delay computation.

3.1 Overview

The utruchirp graphical user interface consists of a single window containing three main sections. At the top, there is a row of tabs that control what is shown in the middle section of the window. This second section, that takes up the largest part of the window is a scrollable view that can accommodate content of any size. Finally, the third section at the lower end of the window contains the monitoring controls, that are always visible as monitoring could be used at any time. In Figure 1, a screenshot of the utruchirp window is shown with the tabs at the top, with *INPUTS & OUTPUTS* selected, the middle section showing the corresponding input and output

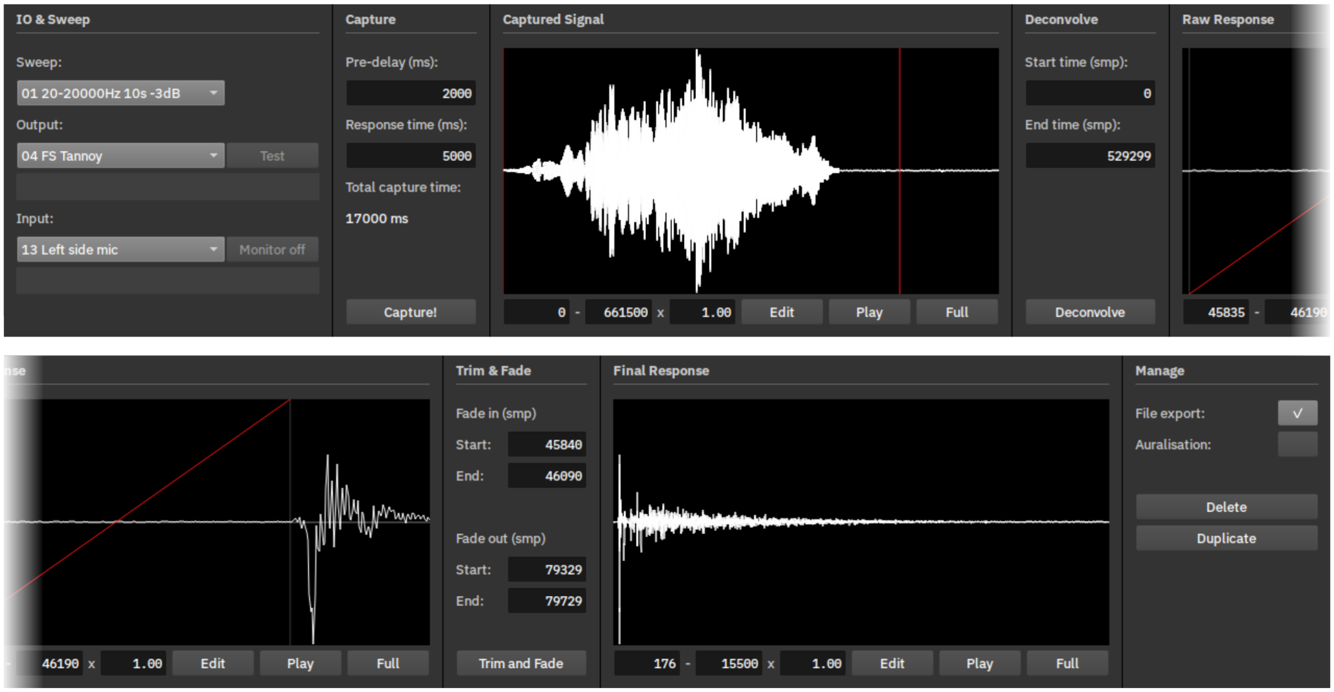

Figure 2: A row from the Measurement tab in utruchirp, here broken up into two parts for readability.

controls, and finally the monitoring section, labelled *MONITOR*, at the very lowest part of the window.

3.1.1 Inputs & Outputs. On this tab, a static list of all the inputs and outputs available to the SuperCollider environment are presented. For a room acoustic measurement, the inputs and outputs would represent microphones and loudspeakers respectively. The inputs can be monitored, named, and comments can be attached to them if there is some special circumstance or piece of information that should be logged and kept with the measurement, for example where a microphone was placed . The outputs are presented in a similar fashion, with testing instead of monitoring. Pressing a test button on an output row sends a test signal to that output. The test signal type and level can be set globally for all outputs. By patching a cable from one physical output directly to a physical input on the audio device used, a round trip delay measurement can be performed, the result of which will be useful at a later stage of the process. Finally, by selecting the inputs and outputs that are to be measured and pressing the *Generate Measurements* button, utruchirp will automatically fill the third tab with enough empty measurement rows to measure all selected combinations, with the correct inputs and outputs set for each row.

3.1.2 Sweeps. The Sweeps tab lists all sweeps that can be used as an excitation signal when measuring. Any number of different sweeps can be added if different parts of the measurement need different sweep characteristics, e.g. duration, start and end frequency, or signal level. For example, if one of the outputs is connected to a sub woofer, using a 20-20000 Hz sweep commonly used for full range speakers might not be appropriate, and similarly, a small speaker can utilise a different sweep that maximises the time the swept sine spends in a frequency range that is actually translated to the measured acoustics. If a speaker is far away from all microphones, the sweep duration can be lengthened to improve the signal to noise ratio of that measurement.

3.1.3 Measurements. Once the inputs and outputs are configured and all needed sweeps are added, one can start the measuring process. The measurement tab contains a dynamic list of rows that each represent the measurement of one input and output combination. Any number of rows can be added, allowing for any number of varying or repeated measurements. Each row represents a process, from left to right on the screen, where a measurement is performed and post processed. Each measurement has its own settings, but wherever appropriate, any parameter value in the process can be set for all measurements by accessing a master control section that always remains at the top of the list. A cropped screenshot of the application window, showing only one measurement row, can be seen in Figure 2. From left to right, the elements of the process and their function are:

(1) *IO & Sweep* - The input, output, and sweep are selected from the sets which have been setup on the previous tabs. Again, outputs can be tested and inputs monitored to aid in selection.
(2) *Capture* - Pressing the capture button activates the capture process. First, the *Pre-delay* number of milliseconds is waited in silence, allowing e.g. anyone present in the space to either settle into a quiet pose, move to another position, or even

leave the room. Then a playback of the measurement's selected sweep on its output is synchronised with a recording of its selected input. The recording is longer than the sweep duration to allow for the reverb tail to be captured. The added length of this *Response time* duration is set depending on an estimate of the measured systems characteristic. The total capture time is therefore the sum of pre-delay, sweep duration, and response time, for each measurement.

(3) *Captured Signal Display* - There are three signal display elements in the row, that have similar functionality. In this element, the captured signal is displayed. There are two view options, displaying either the waveform or a graph of the RMS of the signal amplitude. The display can be freely zoomed in and out in the Y-dimension and display any range of samples in the X-dimension, from 1 sample to the whole captured signal. Typically, one would check for clipping, and make sure that the measured response has receded below the noise floor before the capture ended. If it hasn't, the recording is too short and should be redone with a longer response time. For detailed inspection, the display can go into full screen mode, and if the need for more advanced editing should arise, the captured signal can be opened in an external editor with the click of a button, and automatically reloaded once saved. The display also contains two markers, showing the position of the two parameters of the next element, the deconvolution.

(4) *Deconvolution* - This element processes the captured signal by convolving it with the time-inverted and equalised sweep used in the signal capture, resulting in the raw response signal. The user only needs to set two parameters, the *Start time* and *End time* that together define what part of the captured signal to use. Typically, the round trip delay measured in the Inputs & Outputs tab is subtracted from the beginning of the captured signal, and any part of the response time where the signal has receded below the noise floor.

(5) *Raw Response Display* - This display element is similar to the Captured Signal Display element, and provides identical functionality. The difference is that instead of two markers it displays four, in two pairs, where each pair is connected with a line ramping up or down, to indicate whether a fade out or a fade in is represented. In Figure 2, the fade in markers are visible in the display.

(6) *Trim & Fade* - The raw response needs to be trimmed, and the resulting trimmed section needs to be faded in and out. The four parameters set the start and end of the fade in and fade out. The positions of the two fade and trim operations are displayed in the previous element, the Raw Response Display.

(7) *Final Response* - The third display element shares identical controls and functionality with the other two, with the difference that no markers are displayed.

(8) *Management* - The final element provides options for including the measurement row in the auralisation on the auralisation tab, as well as for exporting it together with other measurements as a set of audio files and a configuration file compatible with Jconvolver. The element also has buttons for deleting and duplicating the measurement row.

3.1.4 Auralisation. The auralisation tab allows the user to listen to an auralisation through real-time convolution using all or a subset of the measurements available in the current utruchirp session without leaving the program and thereby the measurement process. The tab provides a varied set of test signals that can be sent into the auralisation, or to the outputs listed in the Inputs & Outputs tab. Sending identical test signals into the auralisation and the physical space currently being measured and alternating the monitoring between the two allows for detailed comparison of the measurement simulation and the physical space side by side. This way, the fidelity and nuances of the differences between the two can be explored.

3.1.5 Monitoring. The monitoring options are always displayed at the bottom of the window, regardless of which tab is selected. The monitoring bus is a two channel path that provides mono or stereo monitoring of inputs, signals, or auralisations. Some monitor sources are panned to create a stereo signal, e.g. when monitoring a stereo microphone setup or binaural microphones. Others, like a recorded mono signal, are duplicated to both channels, i.e. panned to the centre. Each of the two channels can be sent to two separate outputs on the audio device used, allowing for instance to use two pairs of headphones to monitor, or to monitor using both headphones and speakers.

The monitoring bus runs through three stages before being distributed to the four outputs described above. First, in the input stage, there is level metering and a level control. The second stage contains a limiter that constricts the signal to a maximum amplitude. The third stage has a mute controller, a level control and a level meter. The rationale for this gain control chain is to allow the user to safely monitor all of the highly dynamic signals present in utruchirp. As an example, using monitoring to find when the reverb tail of a captured signal blends into the noise floor of a signal might necessitate an input gain of 50 dB or more. In such a situation, having a limiter followed by a gain control allows the user to set a monitoring level that will never be exceeded. While this signal processing chain can impart an audible effect on the louder part of a monitored signal from the limiting, it is still much preferred to the option of allowing playback of signals of possibly hearing loss inducing level and/or clipping on the output stage of the audio device.

3.2 Round Trip Delay Estimation

A measured room impulse response always starts with silence as the direct sound from the loudspeaker to the microphone takes time to propagate through the room. The delay of the onset represents the line-of-sight distance between the transducers. But there is also a delay incurred by the measurement system due to AD/DA conversion and various other signal buffering stages. Typical round trip delays range from a few to a few hundred milliseconds, representing a few decimetres to many meters of spatial distortion. In order to compensate for this delay, it has to be estimated in a special step involving an analogue loop back connection between an audio input and output of the measurement setup. There are various ways of estimating the round trip delay yielding different degrees of precision. utruchirp uses a simple step function based approach with a precision of about 1 sample, representing a spatial distortion of less than a centimetre, which is considered sufficient based on our practical experience.

If we ignored this additional delay when modelling the transducer/room system through auralisation, our simulation would not be accurate and the system may behave radically differently, especially when used in feedback processes. Such artificially added delay would not only seemingly enlarge the distance between the transducers, but also render the later parts of the room response containing the reflections incoherent, thus skewing the acoustic portrait of the room geometry.

It is critical to acquire this information when it is available, i.e. in close temporal connection to the actual measurement. As the round trip delay may change according to system parameters that are hard to control and monitor, related to the way the operating system is handling the audio I/O, the information may be lost if not acquired at the right moment, i.e. when configuring the measurement setup. Depending on the operating system, it may be hard or even impossible to exactly reproduce the conditions of measurement at at later time. Without the round trip delay compensation a room impulse response measurement may be worthless.

3.3 Monitoring

Room impulse measurement is an iterated multi-stage process. Typically many transducer pairs are measured sequentially and each measurement includes several stages, whereas subsequent stages depend on the data acquired in previous ones. In each iteration and stage a number of things can go wrong, which is why persistent real-time monitoring of all involved signals is essential for a successful measurement session. Typical measurement errors are due to damaged transducers that start to distort under certain conditions, connections that fail or damaged cables being used, accidental noise interjected during the signal capture, or objects in the room that start to rattle due to resonances with the test signal - to name just a few. Most of those problems can be discovered very easily by listening to the signals in question, which is why auditory monitoring is of central importance. In connection with precise signal analysis and visualisation tools, ceaseless listening to all stages of the measurement process is essential for quality assurance.

Furthermore, this type of sensual immersion into the process and its intermediate results gives important insights into how the particular perspective of the room/transducer system one chooses to adopt through the measurement can be made artistically productive. When the immersion into the acoustics of the room through the measurement process becomes a visceral experience, it may substantially inform aesthetic strategies and artistic processes. In order to make rigorous acoustic measurement productive for artistic practice, it is essential that the measurement process does not remain a black box, but that one may relate to it conceptually and sensually on as many levels as possible.

utruchirp exposes the measurement process in form of a matrix allowing for direct comparisons between different stages of one transducer/room configuration (i.e. horizontally) or across different configurations (i.e.vertically). An auditory exploration of such differences allows for a detailed understanding of the acoustic features of the room under exploration. This way one may perceive the room from many different perspectives, building a detailed picture of the acoustic situation we are dealing with in our imagination. Being able to imagine on the level of the features extracted through the measurement is an important prerequisite to make artistic use of the acquired data and the simulations based upon them.

3.4 Auralisations

The artistic practice motivating and informing the development of utruchirp is based to an important extent on simulations of electroacoustic setups through convolution of measured room impulse responses. For that purpose, the room response of each loudspeaker/microphone combination in a setup is measured individually and independently. In a simulation, the sequentially measured responses are combined linearly to account for their interactions. The simulation of each microphone signal consists in convolving all loudspeaker signals with the respective measured response and summing the results of all those convolutions. Thus, the number of convolutions necessary for a complete simulation is equal to the product of the number of microphones and loudspeakers. The convolution matrix of such a simulation has as many inputs as there are loudspeakers and as many outputs as there are microphones in a setup. In this way the simulation of a setup is interchangeable with the real setup. A simple but very important way of testing the validity of the simulation for a particular artistic application is to compare them side-by-side. utruchirp enables such a comparison through auralisation as described in Section 3.1.4.

In our work, a typical simulation will produce signals to be listened to in an auralisation as well as signals which will serve for the feedback process to be composed. It may also be useful and instructive to listen to the latter, but we will usually be more interested in listening to the (typically binaural) signals representing different locations at which the simulation of a setup can be experienced through headphones. Of course, nothing prevents us from also using the binaural signals in the feedback process, which will be useful if we use the corresponding dummy head microphone in an actual performance. As the frequency-dependent directionality of a microphone has an important impact on the particular dynamics of the feedback processes it is used for, it is interesting to exploit the directionality of a dummy head microphone, which has more detailed structural features than the typical polar patterns of most other microphones.

In the current version of utruchirp the convolution process is realised with Jconvolver. To perform a simulation, the convolution matrix has to be specified in form of a Jconvolver configuration file. utruchirp allows for the automatic generation of configuration files based on a selection of measurements in the current session. This greatly simplifies a quick and reliable setup of an auralisation, which involves the creation of the configuration file, starting a Jconvolver process with this configuration file, and establishing the necessary connections between SuperCollider and Jconvolver via Jack. Besides the built-in auralisation feature, utruchirp also allows to export complete simulation setups consisting of the Jconvolver configuration file and the corresponding sound files holding the room impulse responses needed for the convolutions. Such a simulation setup can then also be used standalone with Jconvolver in combination with any other Jack-enabled sound synthesis and processing tools or with the convolution plugins of the VST multichannel plugin suite *mcfx* developed by Kronlachner [10], that uses configuration files in the same format as Jconvolver.

4 FUTURE DIRECTIONS

In this section, we will point to some future directions and highlight some future features that are to be integrated in upcoming versions of utruchirp. These future directions are shaped by insights extracted from of our artistic work, i.e. the composition and performance of several pieces, as well as technical work and studio experiments.

4.1 Experiential Validation

We expect important new insights to be gained from situations in which a direct comparison between the behaviour of a system in a simulated and real room will be possible. This assumption is based on an informal experiment that has been conducted in the context of the CoS project in 2012. Using a headphone equipped with a switch that detected if a listener was wearing the headphone or not, a situation was created where a sound environment created with a large number of loudspeakers in a big concert hall could be experienced - or a static binaural simulation thereof. Switching between the two with a rapid crossfade was simply possible by putting on or taking off the headphone. Most listeners reported to be unable to distinguish between the real and the rendered sound scene, i.e. they believed the speakers in the room continued to emanate sound when they put on the open headphones, but they were actually switched off. Most also reported of a similarly seamless transition when taking off the headphones again and the speakers were switched on as a consequence. Being able to make this type of experience, i.e. knowing the (possibly small) differences between the simulated and real situation is important to base artistic work on such a simulation-driven approach. In the case of our practice the difference between a real and simulated sound scene may be much more pronounced because of the recursiveness of the feedback processes we use. Being able to experiment with those processes in real rooms and their simulations in parallel will be an important opportunity to further develop our artistic practise and decide, for instance, if a combination of real and simulated spaces would be of interest even in a performance situation. This would allow for approaches impossible in a real acoustic situation (e.g. the decoupling of feedback processes) to take place in parallel in a simulation of the very room shared by the performers and their audience in a concert. utruchirp has been designed to perform the described type of experiments.

4.2 Estimation of Nonlinearities

When Farina introduced the swept-sine technique 20 years ago, he framed it as an approach allowing for both estimating the linear impulse response and the nonlinear distortions of a dynamic system. Most of the existing tools implementing the technique (including ALIKI) ignore the distortion products, which appear as pre-echos to the estimated impulse response itself. It is one of the great features of the technique that it allows to easily distinguish the distortion from the signal of main interest, the linear part of the system response. In the current version of utruchirp, a detailed consideration of the distortion products is not supported, but the software has been conceived with this feature in mind, which is planned to be implemented in a future version.

Being able to model the nonlinearities of a transducer/room system, which most likely are caused mainly by the loudspeaker, would be highly interesting in the context of the artistic practice that motivated the development of utruchirp. When composing feedback processes with electroacoustic systems, the role of nonlinearities is central. For instance, in analogue feedback systems the nonlinear distortions introduced through the saturation of various components in the signal path are responsible for keeping an oscillating system in check by diverting the energy from the main spectral components to higher partials and thus limiting the exponential increase of amplitude. Important musical qualities of feedback systems can therefore be shaped by nonlinearities, so they play an important part in designing feedback systems for artistic purposes.

Being able to estimate the contribution of the distortions stemming from the transducer/room system to a feedback process would allow for a better understanding of the phenomena underlying an artistic use of feedback. If the distortion products quantified in the measurement could be used to model the nonlinearities of the transducer/room system, a more faithful simulation of the system could be attained. This would certainly be useful for studying the behaviour of a system in more detail, but most importantly it would allow to explore the relationship and mutual influence of nonlinearities inherent to the system and those introduced in the compositional process.

4.3 Replacement of Jconvolver

Another feature planned for a future version of utruchirp is the integration of the convolution process as a unit-generator, i.e. a native building block, in SuperCollider. Convolution modules for the SuperCollider server already exist, but they are not as efficient and flexible as Jconvolver, currently used by utruchirp. Using a command line application like Jconvolver adds complexity to the implementation and makes it depend on the Jack audio connection kit used to connect SuperCollider and Jconvolver. Performing the convolution operations in the SuperCollider server would greatly simplify the implementation. It would also allow for more flexibility to implement special features, such as dynamising the simulation by interpolating between different impulse responses. Furthermore, it would increase the cross-platform compatibility, especially with Windows, as Jack would not be needed any more. Another added feature that would increase compatibility would be supporting the *Spatially Oriented Format for Acoustics* (SOFA) format [13, 17] for loading and storing sets of impulse responses.

4.4 Vibrant Listening

When we measure a transducer/room setup we have to measure all microphone/speaker combinations individually and then linearly combine them into a simulation of the whole setup. Besides these electroacoustic subsystems, which we turn into coupled oscillators of some sorts in our pieces, we usually also measure the responses of all loudspeakers for several static listening positions using a dummy head. This allows for a comparison of the different binaural listening perspectives. But one problem remains, which is a serious defect of static binaural rendering for headphones - the limited degree of immersion, the problem that gave StiffNeck its name. There are possibilities to make the binaural rendering more dynamic by using multiple-orientation binaural room impulse responses allowing for a dynamic variable-orientation rendering [21]. This involves to acquire a series of measurements, e.g. binaural responses measured every 15 degrees in the horizontal plane. The special iterated measurement mode needed in this case has been foreseen in the design of utruchirp and will be implemented in a future version. It will add a 3rd dimension (iteration) to the current two-dimensional matrix (measurement/stage). This will include also the integrated control of external devices used to change the orientation of the measured microphone or speaker, such as a rotation table in the case of a binaural measurement.

While there is merit to using technically accurate models for improved rendering of the listening position, there are potentially other approaches to increasing the vibrancy and richness of the listening experience. The goal would be to reduce the singular quality resulting from the completely static listening position and the fact that the simulated acoustics is the result of layered static measurements that, unlike real rooms, do not change over time. To do this, the auralisation section of utruchirp would be augmented with a modulation section where slight and subtle modulations could be applied to mimic the effects of small head movements and other variations that are lost in the perfectly frozen moment that is the foundation for the acoustic simulation. Here, an open ended sound design approach would be attempted, where anything that takes us closer to the experience of listening to the physical space that is being modelled, would be considered valid. This, in turn, again highlights the relationship between the notion of accuracy of measurement, where traditional positivist evaluation is desirable, and the experience-centered artistic valuation of how the simulated acoustics can operate in a creative context. In fact, this direction of inquiry could be extended well beyond the goals of simulating passive presence into active movement through the listening environment, all the way into the creation of acoustic experiences that would be impossible to have in a physical acoustic space.

5 CONCLUSION

utruchirp is a highly specialised tool that manifests clear concerns emerging from an ongoing artistic practice. While these concerns are certainly particular to our work, we would like to open up the tool for collaboration and cross-fertilisation. The future work and features suggested in this paper should both be taken as promises of upcoming possibilities as well as calls for collaboration. Our vision is that we can bring as much subtlety and nuance to utruchirp as we do to our digital musical instruments, making the experience and practice of measurement feel like a performance. The utruchirp source code repository can be found at: utrumque.com/utruchirp.

ACKNOWLEDGMENTS

This work has been supported in the context of the artistic research project *Inter_agency: Composing Sonic Human-Computer Agent Networks* (https://interagency.iem.at) funded by the FWF Austrian Science Fund (PEEK AR 483), and by the NAVET center at the KTH Royal Institute of Technology (https://www.kth.se/navet).

REFERENCES

[1] Fons Adriaensen. 2006. Acoustical impulse response measurement with ALIKI. In *Proceedings of the Linux Audio Conference*. 9–14.
[2] Fons Adriaensen. 2006. Design of a Convolution Engine optimised for Reverb. In *Proceedings of the Linux Audio Conference*. 49–53.
[3] Braxton Boren and Agnieszka Roginska. 2011. Multichannel impulse response measurement in matlab. In *Audio Engineering Society Convention 131*. Audio Engineering Society.
[4] Gerhard Eckel and Martin Rumori. 2014. StiffNeck: The Electroacoustic Music Performance Venue in a Box. In *Proc. International Computer Music Conference*. 542–546.
[5] Gerhard Eckel, Martin Rumori, David Pirrò, and Ramón González-Arroyo. 2012. A Framework for the Choreography of Sound. In *Proc. International Computer Music Conference*. 504–511.
[6] Ludvig Elblaus and Gerhard Eckel. 2020. Acoustic Modelling as a Strategy for Composing Site-Specific Music. In *Proceedings of the 15th International Audio Mostly Conference (AM'20), September 15–17, 2020, Graz, Austria*.
[7] Angelo Farina. 2000. Simultaneous measurement of impulse response and distortion with a swept-sine technique. In *Audio Engineering Society Convention 108*. Audio Engineering Society.
[8] Angelo Farina. 2007. Advancements in impulse response measurements by sine sweeps. In *Audio Engineering Society Convention 122*. Audio Engineering Society.
[9] Florian Hollerweger and Martin Rumori. 2013. Production and Application of Room Impulse Responses for Multichannel Setups using FLOSS Tools. In *Proceedings of the Linux Audio Conference*.
[10] Matthias Kronlachner. 2014. Plug-in suite for mastering the production and playback in surround sound and ambisonics. *Gold-Awarded Contribution to AES Student Design Competition* (2014).
[11] Stéphane Letz, Dominique Fober, Yann Orlarey, and Paul Davis. 2004. Jack Audio Server: MacOSX port and multi-processor version. In *Proceedings of the Linux Audio Conference*.
[12] Stéphane Letz, Yann Orlarey, and Dominique Fober. 2005. Jack audio server for multi-processor machines. In *Proceedings of the Linux Audio Conference*.
[13] Piotr Majdak, Yukio Iwaya, Thibaut Carpentier, Rozenn Nicol, Matthieu Parmentier, Agnieszka Roginska, Yôiti Suzuki, Kankji Watanabe, Hagen Wierstorf, Harald Ziegelwanger, et al. 2013. Spatially oriented format for acoustics: A data exchange format representing head-related transfer functions. In *Audio Engineering Society Convention 134*. Audio Engineering Society.
[14] James McCartney. 1996. SuperCollider: a new real time synthesis language. In *Proc. International Computer Music Conference*. 257–258.
[15] James McCartney. 2002. Rethinking the computer music language: SuperCollider. *Computer Music Journal* 26, 4 (2002), 61–68.
[16] David Murphy and Flaithrí Neff. 2011. Spatial sound for computer games and virtual reality. In *Game sound technology and player interaction: Concepts and developments*. IGI Global, 287–312.
[17] Andrés Pérez-López and Julien De Muynke. 2018. Ambisonics directional room impulse response as a new convention of the spatially oriented format for acoustics. In *Audio Engineering Society Convention 144*. Audio Engineering Society.
[18] Julian Vanasse, Andrea Genovese, and Agnieszka Roginska. 2019. Multichannel impulse response measurements in MATLAB: An update on scanIR. In *Audio Engineering Society Conference: 2019 AES International Conference on Immersive and Interactive Audio*. Audio Engineering Society.
[19] M. Vorländer. 2007. *Auralization: Fundamentals of Acoustics, Modelling, Simulation, Algorithms and Acoustic Virtual Reality*. Springer Berlin Heidelberg.
[20] Scott Wilson, David Cottle, and Nick Collins. 2011. *The SuperCollider Book*. The MIT Press.
[21] Markus Zaunschirm, Matthias Frank, and Franz Zotter. 2020. Binaural Rendering with Measured Room Responses: First-Order Ambisonic Microphone vs. Dummy Head. *Applied Sciences* 10, 5 (2020), 1631.