

Production and Application of Room Impulse Responses for Multichannel Setups using FLOSS Tools

Florian HOLLERWEGER

A – 4020 Linz
Austria
flo@mur.at

Martin RUMORI

University of Music and Performing Arts Graz
Institute of Electronic Music and Acoustics
Inffeldgasse 10 · A – 8010 Graz · Austria
rumori@iem.at

Abstract

We present the outcomes of a series of room impulse response (IR) measurements. We have recorded binaural, Ambisonics-encoded and regular stereo/mono IRs of multichannel loudspeaker arrays in various concert halls in Austria, Northern Ireland, Germany and New Zealand. The resulting IRs and accompanying documentation have been made publicly available on a website for composers to use in the production and documentation of their multichannel pieces. The paper also discusses several custom-written shell scripts and extensions to the *Aliki* and *Jconvolver* software packages, which we have developed for the production of the presented IRs.

Keywords

Impulse response, binaural auralisation, virtual acoustics, convolution reverb

1 Introduction

In an earlier paper [Rumori et al., 2010], we have discussed *binaural room impulse responses* (BRIRs), i.e., impulse responses (IRs) recorded using a dummy head in a ‘real’ room, as opposed to an anechoic chamber. We have described the usefulness of BRIRs for the pre-production and documentation of multichannel electroacoustic compositions. By recording the BRIRs of all loudspeakers in a multichannel electroacoustic concert hall, one can generate binaural mixdowns of pieces composed for and/or performed on that loudspeaker array. When listened to on stereo headphones, such a mixdown preserves not only the hall’s reverberation characteristics, but also the perceived 3D direction of all loudspeakers and phantom sources between them. BRIRs thus combine the advantages of single-source stereo IRs used in standard convolution reverbs with those of head-related IRs usually recorded in anechoic chambers.

Binaural mixdowns created from BRIRs assist the composer with the pre-production of multichannel pieces written for concert halls where access is often restricted and rehearsal

time limited. Moreover, they provide an efficient method for documenting multichannel pieces in a format that can be played back on a simple pair of headphones, but nevertheless conveys the acoustic spatiality intended by the composer. Realtime convolution engines such as *BruteFIR* or *Jconvolver* extend the usefulness of BRIRs towards real-time audio input.

Since the publication of our last paper, we have conducted additional room impulse response recordings, using both binaural and other recording techniques, in a variety of concert halls. One purpose of this paper is to provide an overview of these recordings (cf., section 2) and the resulting IR repositories, which we have made freely available (cf., section 5). We also present some examples of how different artists have applied these impulse responses (cf., section 4).

Free/Libre Open Source Software (FLOSS) packages have played a central role in the production of the IRs that we present, and are also essential for their actual application. The creative use of IRs is an area of audio processing where the limited scope of commercial software packages can easily become apparent. The ability to customise many of the FLOSS tools that we have used was both welcome and necessary to achieve our goals. Another purpose of our paper is to provide an overview of our personal extensions to FLOSS tools such as *Aliki* and *Jconvolver* (cf., section 3).

2 Recording Sessions and Venues

2.1 IEM CUBE and SARC Sonic Lab

The *CUBE* concert hall at the Institute of Electronic Music and Acoustics (IEM) in Graz, Austria is a small concert space equipped with 24 Tannoy System 1200 loudspeakers, which are arranged in a hemisphere above the listener and optimised for Ambisonics playback.

The *Sonic Lab* at the Sonic Arts Research Centre, Queen’s University Belfast, is a large



Figure 1: The *CUBE* at the Institute of Electronic Music and Acoustics in Graz (photograph by courtesy of the Graz University of Music and Performing Arts)



Figure 2: The *Sonic Lab* at the Sonic Arts Research Centre in Belfast (photograph by courtesy of Gawain Morrison)

electroacoustic concert hall whose acoustics can be adapted through movable wall panels. At the time of our recording, the Sonic Lab featured 40 full-range speakers (Meyer Sound and Genelec) and 6 subwoofers.

Details on the IRs that we have recorded in the CUBE and the Sonic Lab can be found in our previous paper [Rumori et al., 2010], so we will not reiterate them here.

2.2 New Zealand School of Music

The Adam Concert Room (ACR) is located at the Kelburn Campus of Te Kōkī New Zealand School of Music (NZSM), a joint venture of Massey University and Victoria University of Wellington. The ACR is a medium-size concert hall which is predominantly constructed from wood and seats up to 200 people.



Figure 3: The Adam Concert Room (ACR) at Te Kōkī New Zealand School of Music in Wellington (photograph by courtesy of Stephen Gibbs)

In October 2011, an IR recording session was conducted in the ACR to conclude a one-trimester seminar on *Spatial Audio*. The final repository includes binaural IRs (recorded by a *Kemar 45BA*) and experimental Ambisonics-encoded IRs (*CoreSound TetraMic*) of a circular eight-channel speaker array (*Genelec 1037B/C*), plus standard stereo IRs (X/Y pair of *Neumann KM184s*) of the same array and for a single centred speaker. The repository is accompanied by extensive documentation of the production process and includes several demo scripts to illustrate the usage of the final IRs.

2.3 Academy of Media Arts Cologne

The auditorium of the Academy of Media Arts Cologne (KHM) is a medium-size, multi-purpose hall used for lectures, presentations, film screenings and concerts, the latter mostly of an experimental kind, featuring electroacoustic and improvised music. The reverb time of the hall is rather short as it is mainly conceived for good speech intelligibility.

In November 2011, a hands-on workshop called *Raumfaltung* on measuring and using impulse responses has been carried out at KHM. During the workshop, a “classic” speaker ring of eight *K+H O 108* has been set up in the auditorium and room impulse responses for each of the speakers were measured using an array of different microphones: two kinds of custom-built dummy heads,¹ an ORTF configuration of two *Schoeps MK 4*, a small AB configuration of two *Schoeps MK 2s*, a *Soundfield ST 250* first-

¹Binaural microphones *Sellmeier Grey* and *Sellmeier Brown*, both constructed by Wilfried Sellmeier



Figure 4: The microphone setup at Academy of Media Arts Cologne (photograph by courtesy of Dirk Specht)

order Ambisonics microphone and a *Behringer ECM8000* measurement microphone as a reference channel above one of the dummy heads.

The impulse responses are suitable for mono or stereo convolution reverb applications as well as for auralisation of up to eight channels or binaurally rendered horizontal Ambisonics.

2.4 MUMUTH Graz

MUMUTH Graz is the House of Music and Music Theatre at the University of Music and Performing Arts Graz (KUG). Its main hall, the *Ligeti Hall*, is a multi-purpose space for theatre and opera production, concerts, dance and electroacoustic music. It features a rig of 33 speakers with a motorised position control, which allows for adjusting the heights, pan and tilt of every speaker individually. This way, speaker configurations can be changed and evaluated rather quickly. Additionally, *Ligeti Hall* is equipped with 64 fixed smaller speakers and eight subwoofers, which belong to a virtual room acoustics system, but which can be also directly fed with arbitrary signals. Further speakers and subwoofers may be hung or installed on stands if needed.

Since fall 2010, the artistic research project *The Choreography of Sound*² is carried out at the Institute of Electronic Music and Acoustics (IEM) at KUG. For the case studies of this project, the *Ligeti Hall* is considered the main instrument. As access times to the hall are quite limited, a binaural auralisation of its acoustic properties became desirable, which turned out to be also quite useful for taking up a differ-

²<http://cos.kug.ac.at>, last retrieved 2013-02-17, see also [Eckel et al., 2012]




Figure 5: Speaker rig and measurement setup at *Ligeti Hall* in MUMUTH Graz

ent perspective on the instrument. As part of the *VirtualMUMUTH* subproject (see section 4.4), many room impulse response measurement sessions were undertaken between February 2011 and August 2012, for several combinations of loudspeaker configurations and microphone (i. e. listening) positions.

In order to achieve a high versatility, the measurements were done using an arrangement of multiple microphones: a *Brüel & Kjær* dummy head with additional *DPA 4060* capsules mounted on its temples, a *Schoeps* spherical stereo microphone, an Ambisonics *Soundfield* mic, a large AB configuration of two *Schoeps MK2s* and a mono reference. Due to the high number of channels and the resulting large amount of data, some of the measurement tools needed to be extended as described in section 3.

3 Tools and Customisations

The abovementioned measurements and their post-production have been carried out entirely using FLOSS tools. The central tool being used for all the venues described here is *Aliki* by Fons Adriaensen [Adriaensen, 2006a]. In larger multichannel settings with many speakers, several microphone channels and lots of measurement iterations like at MUMUTH (cf. section 2.4), a few additions and adaptations to *Aliki* became necessary.

For post-processing and using the measured impulse responses, some further tools have been developed, such as scripts based on *sox* for trimming and fading and a script for generating configurations for *Jconvolver*, a high-performance convolution engine [Adriaensen, 2006b].

3.1 Aliko and around

Aliko is a powerful and free impulse response measurement framework, which implements the swept-sine measurement method [Farina, 2000]. It provides means for generating the sweep signal, playing it back and recording the room’s response, for deconvolving the recorded sweep in order to get the impulse response, some functionality for editing the resulting responses and for exporting them.

3.1.1 Automated speaker channel switching

In a larger setup such as in MUMUTH (cf. section 2.4), it is desirable to carry out multiple iterations, e.g. for several speaker channels in a row, in an automated way. Unfortunately, standard Aliko only offers eight output channels and no means of directly switching them automatically. There is a functionality though for specifying the number of iterations and a trigger command, which is executed once every cycle.

For the automated measurements described here, this trigger has been used to send an OSC message to a small SuperCollider patch, which receives the sweep signal on a single input channel (fixed output on channel 1 of Aliko) via the Jack audio server. Upon an incoming OSC trigger message, the sweep signal is routed sequentially to the corresponding loudspeaker output channels. This way, also more complex channel sequences or groups with non-contiguous channel numbers could be easily realised.

For convenience, the custom version of Aliko (as used in MUMUTH) also contained a *reset command* field, which could be triggered by a button. It was used to send another OSC message to reset the channel routing to the starting point.

3.1.2 More input and output channels

Standard Aliko provides eight recording channels for the microphone signals. In MUMUTH, due to the high number of speaker channels (up to 128) and long sweep times (> 10 sec), the actual net measurement time was already so long that it seemed appropriate to add a few more microphones, as those additional impulse responses come “for free” (time wise) and might be suitable for future applications, even if the focus was on binaural room impulse responses. Thus, a configuration of 13 microphone channels arose, including a dummy head, several different stereo pairs, a Soundfield mic (Ambisonics A-format) and a reference channel.

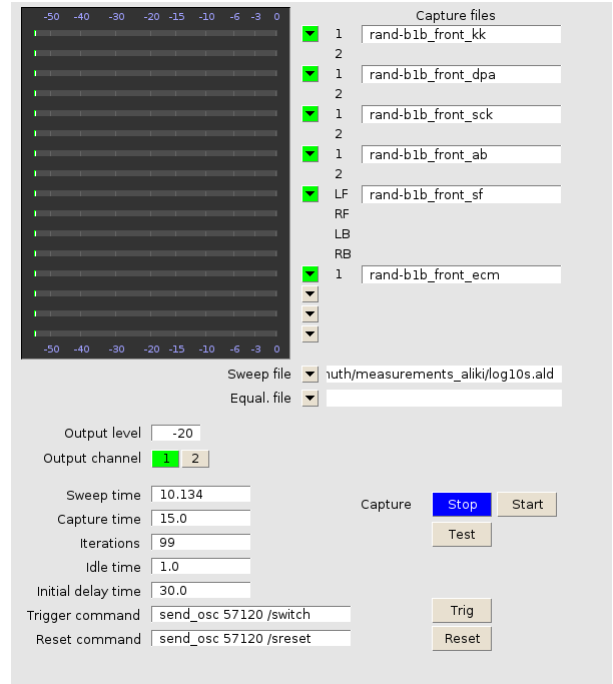


Figure 6: Capture dialog of the extended *Aliko* using 16 input channels, trigger and reset commands and an initial delay

The custom version of Aliko was modified such that the number of audio input channels could be specified as a command line option at startup. In the same manner, also the number of output channels has been parametrised, although this was not needed for the scope of these measurements.

3.1.3 Initial delay

While carrying out a large, automated sequence of iterations, it might be convenient to leave the space being measured in order to minimise the extra noise induced (and to minimise the sweep level in the poor operator’s ears). In the custom version of Aliko, an *initial delay time* may be specified which allows for leaving the space well before the first sweep starts.

Before Aliko was customised for that, sometimes the first measurement (first sweep playback) was regarded a “Nop” and discarded later. This creates problems such as reduced flexibility (fixed sweep length depending on the measurement) and messed-up channel numbers when exporting the impulse responses to single files automatically.

3.1.4 CLI deconvolution and export

Aliko includes the deconvolution stage of the recorded sweeps in order to generate the actual impulse responses, and an export function for

converting from Alikı’s sound file format (`.ald`) to standard formats such as `.wav` or `.aiff`. Alikı’s file format supports multiple sections in a single file, which correspond to the iterations of a sequential measurement. Both the deconvolution and the export may be performed for a single section of a file or for all sections of that file, using the same parameters. The latter option comes in handy when post-processing measurements with a high number of iterations.

However, the measurements reported here, especially those in MUMUTH (cf. section 2.4), not only included many iterations per file but also a large number of those files, due to the high amount of combinations of loudspeaker configurations and microphone positions. Therefore, the tedious and error-prone task of manually loading each file for deconvolution and export was replaced by two custom command-line applications, `aliki-convol` and `aliki-export`, which are basically factored-out versions of the respective built-in functions.

Corresponding to the GUI options in Alikı, `aliki-convol` allows for specifying the start and end times of the deconvolution and whether one or all sections of the file shall be processed. Likewise, `aliki-export` accepts options for the export file format and bit depth, whether a single section, all sections in one file or all sections in separate files shall be exported, the file name base and an offset for file numbering.

3.2 Post-processing impulse responses

All impulse responses taken in the above-mentioned measurement series were not post-processed in Alikı, but through separate tools or scripts. This is mostly due to the high number of responses and the different applications they are targeted at.

Post-processing here basically means trimming the impulse responses, compensating the measurement system latency at the beginning of the file, in some cases also the air travelling time of the direct sound, adding a short fade-in to the beginning, trimming the length of the impulse response and applying a fade-out at the end. Other common tasks are gain adjustments, either by a fixed factor or an individual normalisation, and possibly resampling to a different target sample rate.

3.2.1 `imptrim.sh`

`imptrim.sh` is a shell script used for post-processing the MUMUTH measurements. It is a simple wrapper script around `sox`, which fa-

cilitates the handling of large sets of files with a contiguous number part of a fixed width (e. g. zero-padded) in their filenames. A certain range of numbered files may be processed or all contiguously numbered files found, starting from a given index. For the output files, a different filename pattern and a different starting index for filename numbering may be specified, which enables simple renaming operations for chunks of files at the same time as processing.

`imptrim.sh` allows for separately specifying two portions of the audiofiles to be cut from the beginning, e. g. a compensation portion in samples and another trim in seconds. Further options include the overall length of the resulting impulse response, the length and shape of the fade-in and fade-out and a gain factor. Optionally, a stamp file may be generated in the output directory as a log for the parameters being used for processing.

3.2.2 `remapch.sh`

This is a helper script used for the MUMUTH measurements which actually does not process any audio but only facilitates renaming chunks of numbered impulse response files. It is included here because it might be of general interest for adapting it to other settings.

Measuring and exporting a sequence of speakers whose channel numbers do not start at 1 will result in filenames which do not match the original channel numbers anymore. For example, measuring the MUMUTH sky (channels 65–128) will result in numbered filenames ranging from 1 to 64. This may be avoided by using `aliki-export` and specifying an offset for the output filename (see section 3.1.4).

But when carrying out a large sequence of measurements over a non-contiguous, i. e. “sparse” array of channel numbers, the filename number offset will be different for each loudspeaker block. In MUMUTH for example, when measuring the hemisphere, corner and sky speakers plus two subwoofers in a row, the hemisphere and corner speaker numbers will remain (channels 1–33), the subwoofer channels 42 and 43 will be named 34 and 35, respectively, and the sky channels 65–128 will result in files numbered from 36 to 99.

`remapch.sh` allows for entering such translations for certain regions of channel numbers into a simple database within the script. It may then be used for renaming operations on chunks of response files in both directions, based on the specified configuration (translation).

3.2.3 post_export.sh

The `post_export.sh` script was written specifically for the production of the IRs of the Adam Concert Room at the New Zealand School of Music (cf., section 2.2). It provides an alternative implementation of some of the functions provided by the `imptrim.sh` and `remap.sh` scripts, as well as some additional features.

`post_export.sh` aims at automating the entire IR post-processing after the recorded sweeps have been deconvolved and exported to `.wav`. Similar to the tools presented so far, the idea is to automate the process as much as possible. Thus, all edits are performed with command line tools such as `sox` rather than graphically through the Aliko GUI. Being very much tailored to the specific requirements of the production it was used for, the script is not entirely generic, but hopefully still useful for other IR measurements after according adaptations. The script performs the following processing steps:

Renaming of the IRs exported from Aliko with `mv`, such that they are named according to the respective loudspeaker's direction. For example, the suffix `_nnw` denotes a loudspeaker in north-north-west direction, deliberately (but falsely) assuming that the listener always faces north.

A-to-B format conversion of the experimental Ambisonics-encoded IRs. This step is performed using the `tetrafile` tool by Fons Adriaensen, who also kindly provided the required calibration file for the specific TetraMic (serial № 2099) that was used in the recording.

Latency compensation countering the effects of the recording system's latency, similar to the `imptrim.sh` script.

Trimming the IRs to a pre-defined length with `sox`, similar to `imptrim.sh`.

Fade-in and fade-out with `sox`, similar to `imptrim.sh`.

Normalisation of the IRs to an arbitrary peak value in dB.³ With the help of some bash logic, this is performed in groups, such that

³This step is not strictly necessary, assuming that a floating point number representation is used for the underlying audio files, and considering that in a binaural mixdown of a multichannel piece, many signals will eventually be superimposed anyway.

for every recorded take, the IRs of all loudspeakers in the multichannel setup are normalised to the same target value. The normalisation is deliberately performed prior to converting the IRs to different sample rates, despite the theoretical possibility of clipping due to intersample peaks (which can in practice be avoided by choosing the target value low enough). If an IR was to be normalised *after* resampling, intersample peaks might translate to slightly varying peak levels for different sample rates. The same IR would then be normalised by a different amount for each sample rate.

Resampling from 96 to 44.1 and 48 kHz, using the `sndfile-resample` tool from the `samplerate-programs` Debian package.

3.3 Jconvolver helpers

The tools described in this section do not actually belong to the measurement or post-processing of impulse responses but to their application. In many use cases of the measurements described here, the free command-line convolution engine *Jconvolver* [Adriaensen, 2006b] has been used. Its highly efficient implementation of the so-called *Gardner convolution* using increasing partition sizes makes it especially suitable for long impulse responses and many channels, while at the same time allowing for low-latency operation.

3.3.1 genjconv.sh and genjconvn.sh

Jconvolver uses a configuration file in OSC-message-style syntax for describing the convolution matrix. The settings include the number of input and output channels, the minimum and maximum partition size, the locations, names and channels of the soundfiles containing the impulse response and a couple of options for each impulse response, such as gain or delay.

For the variety of impulse response sets described in this paper, especially those of different combinable loudspeaker groups in MUMUTH (cf. section 2.4), one should be able to quickly and flexibly generate *Jconvolver* configurations. For this task, the `genjconv.sh` and `genjconvn.sh` scripts have been conceived.

Both scripts write a *Jconvolver* configuration file to stdout based on a list of impulse response files read from stdin. The general usage scheme therefore is:

```
ls ir_*.wav | sort -n | genjconv.sh \  
> irjconv.conf
```

The scripts try to be smart with respect to the input and output channels: By default, the number of input channels of the convolution matrix is assumed to equal the number of input files, while the number of output channels is determined from the number of channels of the first input file. Both parameters may also be specified manually, as well as the number of input file channels used and a channel offset.

`genjconv.sh` maps the input channels of the convolution matrix sequentially to the list of input impulse response files. It is possible to specify exclude regions for both input and output channel numbers of the matrix in order to leave “holes”, i. e. produce a “sparse” convolution matrix.

`genjconvn.sh` parses for numbers in the input impulse response filenames and creates a matrix whose input channel numbers are mapped to the corresponding filenames. Currently, this works for zero-padded three-digit numbers directly in front of the last dot in the filename. This facilitates the generation of “sparse” convolution matrices and ensures that impulse response files for certain speaker groups are always associated with the right input channels, independent of the total number of input files given (assuming that they are named accordingly).

For completeness, both scripts allow for specifying a fixed gain factor for all impulse responses, the base path to be used, the minimum partition size, and the maximum impulse response length.

3.3.2 `jjconvolver.sh`

When using manually written or automatically generated *Jconvolver* configuration files across different systems, a common problem is that the impulse responses reside in a different location than specified in the configuration file as an absolute path. Often, the difference is negligible and is caused only by the different name of the home directory or a few initial path elements, as the deeper structure of a commonly used impulse response database is likely to be the same.

`jjconvolver.sh` tries to solve this problem by simply replacing the `/cd` directive in the *Jconvolver* configuration, generating a temporary file and starting *Jconvolver* using the adapted configuration. It acts as a frontend script to *Jconvolver* and passes through any of its options. Additionally, it accepts a few extra options which either specify a full path to completely replace the existing one in the con-

figuration, or the initial elements of the new path from which the complete path is guessed. This is done by concatenating the given initial path and the existing one, subsequently stripping one hierarchy from the front of the existing path, until a directory with the resulting name is found in the filesystem. Although this is a very simple mechanism, it turned out to be quite effective in many practical use cases.

4 Application examples

4.1 Electroacoustic compositions

In our previous paper [Rumori et al., 2010], we have already discussed some artistic applications of our impulse responses by different composers. Justin Yang, Gary Kendall, and John Moeller have used the IRs recorded in the SARC Sonic Lab for the production and documentation of several multichannel electroacoustic compositions. Dirk Specht and Gerriet K. Sharma have published binaural mixdowns of their multichannel pieces on two CDs, using the IRs recorded in the IEM CUBE.

4.2 Pop Production

Nic McBride has used the binaural room impulse responses from the Adam Concert Room at the New Zealand School of Music to create a headphone mix of eight distinct stems in a pop production entitled *L’addio Scontato*. This example represents an original use of BRIRs in the sense that the final production was never primarily targeted at multichannel playback – one cannot really speak of a stereo ‘reduction’ in this context. According to Nic, the BRIRs allowed her to achieve a natural-sounding spatial separation in the final binaural mix, which has been released online [McBride, 2012].

4.3 Audio Augmented Environments

Binaural impulse responses recorded in several of the described measurement sessions have been used by Martin Rumori in sound installations which form different kinds of audio-augmented environments. They include navigable virtual environments incorporating a headphone tracking system, such as *Parisflâneur*,⁴ or mixed-reality installations, where the virtual layer blends with the real auditive surrounding of the listener, using open-loop headphones, such as in *ruhrprotokolle*.⁵

⁴<http://www.rumori.at/do/parisflaneur>

⁵<http://www.ruhrprotokolle.de>

4.4 VirtualMUMUTH

Within the research project *The Choreography of Sound* (see section 2.4), the impulse responses measured in MUMUTH's *Ligeti Hall* are used for an auralisation of the acoustic properties of the space. The *Jconvolver* helpers described in section 3.3 have been integrated with the *CoS* software framework conceived in SuperCollider, which will be published along with the project's documentation in the second half of 2013. The auralisation is completed by a visual three-dimensional navigable model of *Ligeti Hall*, realised in Blender and also integrated with the software library.

Participants of the *CoS* workshop at *Impuls Academy 2013*⁶ in February 2013 and artists invited to the *Mind the Gap!* symposium in MUMUTH in March 2013⁷ prepared their contributions using the *VirtualMUMUTH* standalone application.

5 Availability

The impulse responses from all recording sessions described in section 2, including the patches, scripts and helpers introduced in section 3 as well as the *VirtualMUMUTH* auralisation tool mentioned in section 4.4, have been made freely available at <http://irdb.kug.ac.at>.

6 Conclusion

Impulse responses of multichannel loudspeaker arrays in real concert halls provide many useful applications for music production and the documentation of electroacoustic compositions. FLOSS tools continue to play an essential role in the recording and usage of such IRs, due to the customisability that these tools offer, and the current limitations of commercial software packages.

We are looking forward to feedback from music producers and composers who would like to use the presented impulse responses for their own work, hopefully in unusual and original ways that we have not envisioned here.

7 Acknowledgements

The measurements at MUMUTH Graz were carried out within the *Choreography of Sound* artistic research project initiated by Gerhard Eckel and Ramón González-Arroyo, funded by

the Austrian Science Fund (FWF) under the project code PEEK AR41.

Fons Adriaensen has provided much-appreciated support for his excellent *Aliki*, *Jconvolver* and *Tetrafile* packages.

We would also like to acknowledge the contributions of Bridget Johnson, Jason Post, Stuart Macann, Roy Carr, Dugal McKinnon and Mark McGann at the New Zealand School of Music, and of Mark Poletti at Industrial Research Ltd. in Lower Hutt.

At the Sonic Arts Research Centre in Belfast, we would like to thank Andrés Cabrera, Gary Kendall, John Moeller, Justin Yang and Chris Corrigan.

Thanks to Anthony Moore and Dirk Specht for organising the measurement workshop at the Academy of Media Arts Cologne.

References

- Fons Adriaensen. 2006a. Acoustical impulse response measurement with ALIKI. In *Proceedings of the International Linux Audio Conference 2006*, Karlsruhe.
- Fons Adriaensen. 2006b. Design of a convolution engine optimised for reverb. In *Proceedings of the International Linux Audio Conference 2006*, Karlsruhe.
- Gerhard Eckel, Martin Rumori, David Pirrò, and Ramón González-Arroyo. 2012. A framework for the Choreography of Sound. In *Proceedings of the International Computer Music Conference 2012*, Ljubljana.
- Angelo Farina. 2000. Simultaneous measurement of impulse response and distortion with a swept-sine technique. In *Audio Engineering Society Preprint 5093*.
- Nic McBride. 2012. L'addio scontato. <http://soundcloud.com/amynicmcbride/1-addio-scontato-elyssa-vulpes>. Produced by Nic McBride, written and performed by Elyssa Vulpes. Last retrieved on 12 February 2013.
- Martin Rumori, Florian Hollerweger, and Andrés Cabrera. 2010. Binaural room impulse responses for composition, documentation, virtual acoustics and audio augmented environments. In *Proceedings of the 26th VDT International Convention*, pages 670–679, Leipzig.

⁶<http://www.impuls.cc>

⁷<http://www.researchcatalogue.net/view/33841/37723/1022/567>