

General-purpose Ambisonic playback systems for electroacoustic concerts - a practical approach

Jörn Nettingsmeier

Freelance audio and event engineer,
Lortzingstr. 11, 45128 Essen, Germany
nettings@stackingdwarves.net

ABSTRACT

Concerts of electro-acoustic music regularly feature works written for various speaker layouts. Except in the most luxurious of circumstances, this implies compromises in placement, frequent interruptions of the concert experience to relocate speakers, and/or error-prone equipment rewiring or reconfiguration during the concert.

To overcome this, an Ambisonic higher-order playback system can be used to create virtual speakers at any position as mandated by the compositions. As a bonus, the performer can then be given realtime control of the source positions in addition to their levels, increasing the creative freedom of live sound diffusion.

Deployments at LAC 2009 and the 2009 DEGEM concert at musikFabrik in Cologne have yielded very good results and were met with general approval. Additionally, two informal listening tests with a group of film sound artists and electro-acoustic composers were conducted to gather expert opinions on advantages and shortcomings of the proposed virtual speaker system.

1. INTRODUCTION

This paper describes a practical Ambisonic concert system, as well as a setup procedure suitable for live situations. The system consists of a Linux-based rendering machine, and otherwise standard sound reinforcement equipment that should be available at electro-acoustic music facilities or is readily obtained from rental companies.

All software parts of the production toolchain are free and open-source. Hence, they do not incur licensing costs, are reasonably easy to port to and interface with existing Mac OS X, Solaris and (with certain limitations) Microsoft Windows environments [Sle10], and are easily customized to accommodate special requirements that frequently arise in any experimental arts context.

The sound engineer will appreciate the ease of creating different speaker layouts or trying small layout variations to better convey the composer's intentions in a given venue, and the ability to compensate for non-optimal speaker positioning due to space constraints, escape routes and other venue safety requirements.

Concert curators will like the added flexibility in accommodating non-standard spatial configurations (whose variety grows even more quickly when height reproduction is desired) and the option of commissioning or accepting native Ambisonic compositions without extra cost or error-prone pre-rendering attempts by composers. With changeover times and distracting stage work minimized, it becomes much easier to create a coherent and focused concert experience.

2. OVERVIEW

A minimal Ambisonic live playback rig consists of a multichannel playback source with appropriate encoders, a decoder, digital-to-analogue converters and a number of speakers and amplification.

For pre-produced tape works, the source, the encoders (i.e. the panners) and the decoder will usually run on a single PC. If the work features live electronics or sound is rendered in real time on a machine provided by the artist, it will be necessary to accommodate external audio sources. By providing analog line-ins as well as multichannel digital inputs, all situations should be covered.

While the decoding machine can drive the converters and amps directly, it is highly desirable to have a physical master volume control outside the computer: for convenience, and as an emergency breaker in case something goes horribly wrong (as happens with complex digital systems from time to time).

A digital mixer with digital I/O is the obvious choice. It also handles analog signals from an announcer's microphone or traditional instruments, if necessary.

For flexibility, the mixer's inputs should be routable through the encoding machine and back into the mixer, so that external signals fed into the mixer can benefit from Ambisonic panning as well. In such a setup, it is important to consider (and minimize) the system latency¹.

2.1. External signal transmission

Currently, there are two suitable multichannel digital interface standards that are well-supported by free software.

ADAT lightpipe connections provide eight channels of 24 bit audio at 48 kHz over a cheap optical link using polymer fibre with plastic toslink connectors. Their maximum transmission distance is specified as 10 metres; under optimum conditions, 20 m are possible. ADAT has become a commodity standard that is available on many consumer and professional computer interfaces, many of which have production-quality Linux drivers. Most live mixing desks support ADAT natively or through optional interface cards. The channel count is usually limited by the number of available extension bays. 32 inputs and outputs (i.e. four ADAT pairs) are a common upper bound, which is suitable for Ambisonics up to fourth order plus room for subwoofers and auxiliary feeds.

¹The ffado.org website provides a good starting point for latency tuning. The focus is on firewire devices, but most of the tricks are applicable to all audio systems: <http://subversion.ffado.org/wiki/LatencyTuning>

MADI (Multichannel Audio Digital Interface) connections carry 56 channels of 24bit audio at 48kHz (or 64 if the optional varispeed capability is not needed) over either coaxial 75 ohm cables terminating in BNC connectors (which are cheap and reliable), or optical fibre using ST plugs (slightly more expensive, but an ace in the hole over long distances where hum currents might be an issue). Default MADI sockets are rare on mixing desks, but most vendors offer optional expansion cards. On the computer side, as of this writing, the choice is currently limited to the RME MADiface (available in PCIe and Express-Card flavours), for which complete and reliable Linux support is available.

ADAT-to-MADI bridges and vice versa are available from several vendors for more complex setups.

Both ADAT and MADI can accommodate double or quad sample rates by demultiplexing the signals onto two or four physical channels, at the cost of reducing the number of transmission channels accordingly. It should be noted that the only tangible benefit of higher sample rates in a live situation is a slightly reduced latency (since most devices have a constant number of samples of buffering, regardless of the rate), but at a significant cost in bandwidth, CPU and storage.

2.2. Software components and internal signal flow

The decoding PC should be equipped with at least two ADAT I/O connectors, i.e. 16 channels in each direction. Tape pieces can be played back from the PC itself, and live electronics, microphones or instrument signals will be coming in via the ADAT inputs.

Inside the PC, the audio signals are handled by the **JACK Audio Connection Kit** [Dav10], a low-latency sound server that allows the user to route audio between different applications and have them co-operate in real time.

In the use case at hand, these are:

- a *playback engine* for tape pieces,
- a *virtual mixer/signal router* for manipulation and to host the panning plugins, and
- an *Ambisonic decoder* to generate the speaker feeds.

This author prefers **Ardour** [Dav10-2], a very flexible, free digital audio workstation, as the signal handling hub inside the PC, since it can double as a playback engine for complex multichannel pieces and offers parameter automation. If only routing and panning are desired, a more lightweight mixer application such as the Non-Mixer [Lil10] could be considered.

All incoming signals are patched into Ardour buses (or tracks, if recording capability is necessary), tape pieces are imported as single-channel audio regions into mono Ardour tracks, and a master bus suitable for the desired target format is created. In the case of 3rd-order periphonic playback, it is 16 channels wide.

Note that Ardour itself is agnostic to surround formats – it is only the panner that cares about (and defines) the meaning of each channel in the master bus.

The default panners in all tracks and buses must be bypassed. Instead, an Ambisonic panning plugin is inserted as the last post-fader plugin of the channel strip. Panners up to third order are freely available as part of the **AMB plugin set** [Adr10]. An extension to higher orders is trivial.

For tape pieces, the panners can be set in advance to produce the required virtual speaker positions, either statically or with automation, if in-flight adjustments are required.

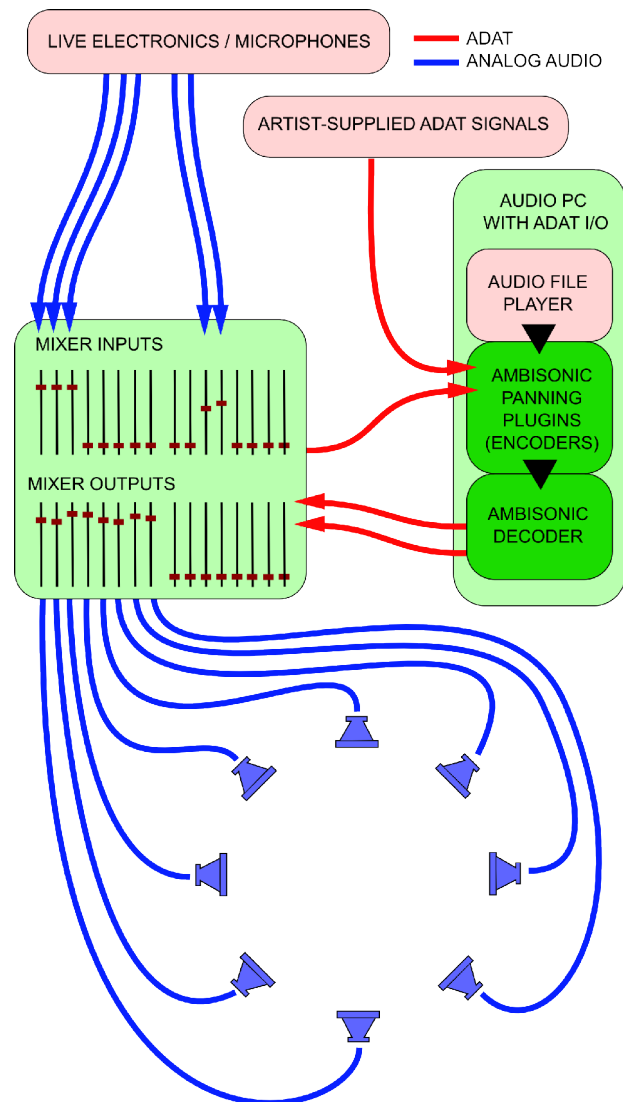


Illustration 1: Signal flow overview of an Ambisonic concert playback system, using a digital mixer. Mixer output faders used for speaker gain compensation

If necessary, the performer can be given real time control of rotation, azimuths and elevations via MIDI controllers. The mapping of parameters in Ardour is as simple as a shift+control middle-click on a fader widget to initiate MIDI learn mode, followed by a quick wiggle of the desired hardware controller, which is now bound to the widget.

Optionally, an Ambisonic rotator can be inserted into the master bus to correct for angular offsets of the rig (if, for example, a two-in-front octagonal preset is being used for an octagon with a centre speaker), or as an additional control parameter for creative live sound diffusion.

From Ardour's master bus, the signal is then fed to an appropriate JACK-aware decoding software. The author recommends **AmbDec** [Adr10-2], a state-of-the-art decoder capable of up to 3rd-order horizontal and first-order vertical playback. It sup-

ports dual-band decoding and optional but highly recommended near-field compensation.

After decoding, the signals can be patched to the ADAT outs, or optionally run through a convolver for additional FIR filtering, and then back into the physical mixing desk, which in turn feeds the speakers.

3. EQUIPMENT REQUIREMENTS AND SETUP

As noted above, a silent PC with sufficient ADAT I/O is essential. Multi-core CPUs are convenient to ensure a responsive user interface when audio computation load is high, but they are not mandatory. Anything like a 1.0 GHz Pentium or better should be more than adequate for simple playback and decoding.

Budget permitting, a digital mixer is highly recommendable. It should have as many analogue outs as there are speakers to drive.

Next, decide on the make and number of speakers. In Ambisonic systems, all speakers will contribute to any one virtual source to some extent, giving you a slightly higher SPL efficiency over the entire listening area. On the other hand, electronic music does require rather more headroom than your average mastered pop material. For a small audience (50-80 people), 8 or more large active near-field monitors might be adequate in a medium-sized room if supported by subwoofers, but for anything larger, P.A.-grade material should be used to avoid sonic compromises and burnt drivers. All speakers and amplification should be identical, to ensure consistent phase response.

Six speakers will give you second-order horizontal coverage with a usable listening area of perhaps one-third the radius of the speaker circle. Third-order operation is recommended for an extended listening area of one-half to two thirds radius, with improved source sharpness that is closer to (but not quite on par with) discrete speaker panning. The minimum number of speakers for 3rd-order horizontal is eight.

3.1. Selection of a suitable base layout

Assuming that the reader does not have the means of computing her/his own decoding coefficients (much like this author), it is highly advisable to select a speaker layout that can be derived from one of the presets shipped with AmbDec. As of this writing, higher-order capable presets include a number of regular polygons, and regular polyhedra as well as stacked rings for full 3D reproduction².

Generally, the angles of incidence should be kept as dictated by the preset. The speaker distances can then be varied within reasonable limits to accommodate the venue, as they only require delay and gain compensation without affecting the decoding coefficients as such.

Extreme distance differences should be avoided, because the maximum sound level will be bounded by what the farthest speaker can safely deliver at the sweet spot. Moreover, distant speakers will have more room reverb at the listening position,

²For special setups such as hemispheres, the author has found that people who are researching decoders are generally very interested in getting field testing and user feedback. Which means you might get a custom-tailored matrix from a friendly Ambisonics research facility near you if you ask nicely. The AmbDec author has also made a standing offer to try his skills on any interesting speaker layout you might care to throw at him.

which will reduce the sharpness of sources in that direction, create uneven tone color across the rig and might cause closer speakers to dominate the directional perception.

In practice, the usual starting point will be a circle, which can be elongated to an ellipse or even made into a rectangle if necessary.

3.2. Speaker positioning

The first step is to mark a point of reference on the floor in the centre of the listening area. This will be the "sweet spot" to which the setup is being calibrated.

Now set up the speakers at the correct angles (minor deviations are acceptable). A laser angle gauge will speed up this process considerably - adequate specimens can sometimes be obtained cheaply at your local do-it-yourself store.

3.3. Distance measurement

Mount a laser range finder at ear height on a tripod over the point of reference. Then measure the distance to each of your speakers precisely (to within 1-3 centimetres). Reproduction in the sweet spot breaks down as the error approaches $\lambda/2$ (around 6 kHz at 3 cm). Outside the strict sweet spot and in a diffuse field, the distinction is largely academic, but it is suggested that this precision be maintained to get predictable results.

Aim at the tweeter or some other easily identifiable flat surface, but avoid shooting at the front cloth or mesh grilles, as these may introduce errors depending on whether your beam lands on the surface or goes through the material. If there is a bass port behind your measuring spot, the error may be substantial. Similarly, avoid cheap range finders that do the actual measurement with ultrasound, since you can never be sure what you are measuring. To mislead users, those cheap units usually feature a laser as well, but it is only used for aiming, and the actual measurement beam is much wider and more diffuse.

3.4. Distance compensation and level adjustment

Now edit the AmbDec preset file you have based your speaker setup on. Leave the angle values alone (they are only there for clarity - changing them will not automatically adjust the coefficients), but do enter your measurement results into the speaker distance column. These values govern the distance correction functions, which you must activate in the setup dialog: be sure to tick both "delay compensation" and "near field compensation". You can also activate gain compensation based on distance, but that assumes that all your speakers are precisely level-matched to begin with. As this is usually not the case, leave that option off and proceed as follows:

On your point of reference, set up a good SPL metre (pointing straight upwards, to avoid measurement errors due to shadowing effects, at least in the horizontal plane) and calibrate all speakers to within +/- 0.1 dB using pink noise. The calibration gains can be applied in the external mixer or, if available, in the software mixer of your audio interface. If neither of these options is available, use an additional Ardour bus for each speaker send.

As an alternative calibration method, use an omni-directional measurement microphone. A cheap one, like the Behringer ECM 8000, will do fine. Again, point it straight upward or downward. Install the JAPA realtime analyzer [Adr10-3] on your rendering machine and feed the microphone signal into it.

Conveniently, JAPA provides a pink noise test signal, which you can now route to each speaker in turn, using AmbDec's external test signal input. Use the most distant speaker as your reference, bring it up to the desired level, and make sure it is not clipping or limiting. Make a snapshot of the steady-state curve in noise mode (i.e. with slowest response time) with the “-> X” button once the system has stabilized. Keep it on the display by selecting the “X” curve in the memory section. Now iterate over your speakers, adjusting each gain until the resulting measurements match the reference curve as closely as possible. We are only interested in overall loudness, so ignore any minor deviations in frequency response.

A crude level calibration can also be accomplished with a simple metre and a 1 kHz test tone, but then you are susceptible to large adjustment errors if something funny happens to be going on in that narrow test band, like destructive interference with a floor reflection.

3.5. Aside: speaker equalisation

You should avoid 30-band EQs per output to compensate for acoustic deficiencies of room and speakers. The phase deviations introduced by filter banks at different settings would severely degrade the recreated sound field.

Hence, the best approach is to either do without equalisation altogether, or to ensure matching phase responses in all channels. It is a good (and time-saving) compromise to use the same EQ curve (and algorithm!) for all speakers, which is easily achieved by ganging the EQs on a physical (digital) mixer, or by inserting a mono EQ plugin into Ardour's master bus, which will be replicated automatically for all channels. In theory, it is possible to use custom phase-matched FIR filters for each speaker³. However, the benefit will be marginal in most cases, at the cost of additional latency and a setup time and effort that is prohibitive in all but permanent installations.

3.6. Adding optional subwoofers

Depending on the number of woofers available, different optimised driving signals can be used.

In the case of just one woofer, the obvious choice is to feed it the W channel and to place the speaker it in the centre front (where any “residual” localisation cues will be the least distracting). If two are available, they can either be placed centre front and back, or left and right, and be driven with W +/- X, or W +/- Y, respectively. In the ideal case of four woofers, a standard first-order square decoder can be used. If eight units (and safe rigging anchors) are available, place them in a cube and drive them in first-order periphonic. Higher orders have no advantages at low frequencies.

The user should experiment with different ratios of W to the directional components. More W will increase the efficiency and overall “boom” that the system can deliver. If emphasis is placed on the directional components, a nice dry trouser-flapping effect can be observed, without the obnoxious pressure on the ears that is often associated with deep bass in enclosed spaces. Obviously, the latter option reduces the maximum sound pressure that can be delivered.

To derive the signals, you have two options: either run two instances of AmbDec in parallel, one for the tops and one for the subs, or create a custom configuration that includes both. The

latter is generally more convenient to use, while the former is less work to set up.

Subwoofers should be aligned and calibrated using the procedure described for the tops. With respect to equalisation, the same rules apply. For band separation, use low- and highpass plugins⁴.

4. TESTING AND FINE-TUNING

After the system has been set up and calibrated, it is important to check for defects with well-known program material.

As a first step, a mono signal should be routed to each adjacent pair of speakers, one after the other, to check for polarity errors. If the polarity is correct, a stable phantom image should appear in the middle of the speaker pair under test.

Next, pan the same signal slowly across the entire rig and ensure there are no major changes in loudness.

After any problems have been corrected, send an ambisonic test signal with sustained HF content over the system. Depending on the room acoustics, you will experience irritating phasing effects when moving around near the centre spot. They will be quite pronounced in very dry spaces and inconspicuous in reverberant rooms. Decide whether that phasiness is an issue with the program material at hand – actual audibility will depend on HF content and on whether the audience is seated or invited to move around.

You can trade in some of the phasiness in the centre for a slight loss of localisation precision by “de-tuning” your carefully measured delays. As a quick A/B check, temporarily disable delay compensation in AmbDec to hear the effect.

Next, experiment with the crossover frequency between the LF and Mid/HF decoder⁵. Start with the default (which is always very conservative), and move it upwards a hundred hertz or so, to see if localisation clarity improves or degrades.

If you have the chance, try and compare a discrete stereo signal sent to two speakers with the same signal panned in Ambisonics, and listen for coloration. Likely, you will observe a slight damping in the treble range, which can be corrected with one or two dB boost at 5-6 kHz with a bandwidth of two octaves.

Sometimes, venue constraints such as escape routes may force you to move a speaker away from its optimum angle. Sources will now be drawn to where two speakers are lumped together, and a “hole” will open up on the other side. In this case, try to raise the level of the lone speaker gently, and bring down the level of the other two speakers accordingly. Test the effect with a pink noise signal panned across the problematic area, and aim for constant loudness rather than smooth movement.

5. LISTENING TESTS

Composers performing their works on a 12-speaker 3rd-order horizontal Ambisonic system have reported a very pleasing sense of space and envelopment, free from timbral defects or obvious imaging deficiencies even at the peripheral areas of the auditorium. However, it should be noted that due to time and equipment limitations, no A/B comparison with traditional discrete loudspeaker replay could be offered to them.

To corroborate the author's hypothesis that virtual Ambisonic sources are indeed a workable approach to multichannel play-

³Another free software package is available for this job: DRC-FIR [Sbr09]. [Net08] shows how DRC can be applied to an Ambisonic listening rig.

⁴6 and 12dB/oct variants are available as part of the CALF plugin set [Fol10].

⁵For a detailed discussion of dual-band decoding, see [Ger74].

back, two informal listening tests with different target groups were scheduled, this time including A/B comparisons.

5.1. Film sound people

The first test session was conducted at the Kunsthochschule für Medien in Cologne, in a Dolby-certified film mixing room with very dry acoustics. The audience consisted of an experienced film mixing engineer, a film composer, two students of media arts, a film projectionist and a media arts Ph.D. student. Due to the limited number and at the same time high expertise of the participants, it was decided not to attempt a statistically valid formalized test, but to collect opinions and discuss impressions instead.

Two rigs of identical loudspeakers (K+H O108TV 2-way 8" systems) were set up for direct comparison: an Ambisonic octagon driven in third order, and an ITU 5.0 setup. Both rigs shared a common center speaker. For the test, 5.0 content (both film and music) was reproduced either natively or as 3rd-order virtual sources on the Ambisonic rig. Obviously, the Ambisonic system could not be expected to perform better than the ITU rig, as it had to suffer from the inherent limitations of the 5.0 source material, and then added its own problems. Rather, the goal was to find out whether Ambisonic reproduction could be considered a workable compromise, and to learn more about the perception of its shortcomings.

Participants were encouraged to move around while listening, explore the usable listening area and watch out for position-dependent artifacts.

In the first listening phase (which featured a short animated film), the most strongly evident problem of the rig under test was the tendency of sources to move along with the listener, i.e. to maintain relative direction but not absolute position. This was unanimously deemed unacceptable for the centre channel in film sound reproduction.

Listeners also reported a much wider frontal sound stage, which some found pleasant, and others criticized as "loss of frontal focus".

All participants expressed a clear preference for native reproduction in the film sound use case. The mixing engineer remarked that she could imagine working with an Ambisonic system, but that the increased stage width would have affected her mixing decisions.

In a second listening phase, various 5.0 music snippets were presented: a Händel aria, a piece for organ and percussion, some lounge jazz and an electric reggae track.

The initial reaction to Ambisonic playback was positive – the listeners reported better envelopment and pleasant ambience reproduction. One participant perceived an irritating amount of phasiness and jumping sources depending on pitch (specifically that transient attacks would seem to come from a different direction than the sustained sound). In direct comparison, two test subjects reported audible coloration.

Half of the participants still expressed a clear preference for 5.0, while the other half was undecided. General consensus was that the Ambisonic system was far better suited to music reproduction, where its characteristics would be less obtrusive or even beneficial.

A couple of tentative conclusions can be drawn:

1. **For cinema use, absolute localisation of the center channel is mandatory.** This requirement cannot be met by a third-order Ambisonic system. In a subsequent test run, the centre

channel was routed around the Ambisonic encoder and directly to the front speaker, with a gain boost of 4dB to maintain balance (the exact amount will likely vary with the number of speakers and Ambisonic order). This hybrid setup was found acceptable. Since that speaker remained part of the Ambisonic rig, it now carried both the direct C signal and components of L/R/SL/SR. The test content did not show obvious artefacts, but further experiments with LCR-panned material will be necessary to check this method for phasing or comb filter effects.

2. **Focus and stability of localisation were considered critical**, while holes in the side and rear sound stages or distorted ambience reproduction or localisable speakers were not. Even when listening to native B-format recordings, the test subjects did not express much appreciation for the advantages of Ambisonic reproduction. Instead, they reported localisation ambiguities and coloration (which were clearly present, but not particularly disturbing to this author, compared to the benefits). This experience suggests that people who have had prior exposure to Ambisonics have learned to "decode" a great amount of spatial detail from Ambisonic playback and will usually prefer it, while subjects without such listening experience tend to be irritated by ambiguities and diffuseness. It might well be that Ambisonic listening takes some getting used to, and that the listeners' verdict might improve over time. But then it also implies that Ambi "professionals" tend to overestimate the impact on (and quality perceived by) "laypersons", i.e. casual listeners.

3. The **widening of the listening area** during Ambisonic playback postulated by the author was **not substantiated** by listener remarks. Interestingly, the collapsing of the image into a side or rear speaker during 5.0 playback at peripheral positions was not commented on either. Apparently it was considered natural and not perceived as a problem.

4. The **dry acoustics of typical cinema rooms will make phasing artifacts very evident.** Additional measures to reduce phasiness (such as de-correlators for the HF band) should be explored under such circumstances.

5.2. Electro-acoustic musicians

A second listening test was conducted at the Institut für Computermusik und Elektronische Medien (ICEM) at Folkwang Hochschule Essen. During a 2-day workshop, one sound engineer, five students and one professor of electro-acoustic composition experimented with an octagonal setup of Apogee AE-8 15/2" speakers in a moderately ambient room. On the first day, the students set up their own quadrophonic compositions for playback and evaluated both native and Ambisonic renderings. On the second day, a short survey was taken.

For this survey, the author had selected seven excerpts, five from student's works and two from other compositions. Each excerpt was first played back natively, then via Ambisonics. Participants were encouraged to move around to be able to assess the usable listening area. After some time for note-taking, each excerpt was repeated, again on both systems. The subjects were asked to compare the Ambisonic playback to the discrete reference, specifically coloration, stability of localisation, angular fidelity, area of very good resp. usable reproduction, sound of phantom sources, smoothness of movement, ambience, envelopment, source distance and size, personal preference, and appropriateness of the Ambisonic reproduction for the artistic work.

Subjects reported **minor shortcomings throughout, but they were mostly deemed inobtrusive.** Only for one piece was the Ambisonic system considered unsuitable: it treated the speak-

ers as individual sources, without any correlated information or intent for phantom imaging. Here, any diffuseness or spatial widening was clearly detrimental.

On average, the subjects expressed **“no preference”** with a slight overall tendency towards native reproduction, and assessed the Ambisonic system to be **“very usable”** for conveying the artistic intention. However, **individual judgement deviated considerably** – in the most extreme case, one person expressed a strong preference for Ambisonics while another considered it clearly unusable for the same excerpt.

Again, the **subjects did not share the author's impression of a slightly enlarged sweet spot**. On average, the listening area for “good” imaging was found to be 0.5 times the array radius for Ambisonics, and 0.6 for discrete; “usable” imaging was perceived up to 0.7 and 0.8, respectively.

In comments it became evident that **subjects wanted to be able to pinpoint speaker positions** – much to the author's surprise, the ability to do so was considered a matter of reproduction fidelity.

After the survey, some music DVDs were played back in both modes, which provided further interesting insights. One specimen in particular, a recent live production of the Eagles' “Hotel California” showed severe comb filtering of frontal sources, to the point where the result was unusable. Detailed inspection showed negative correlation between the C and L/R channels, a production trick that widens frontal sources over discrete speakers, but lets Ambisonics fall flat on its face. A Pink Floyd Live DVD had the main instruments mixed to L/R exclusively, without significant centre channel content. When switched to Ambisonic playback, the intentionally airy and diffuse sound stage would gel into a very focused centre image, certainly “correct” but very audibly different.

It could be argued that these highly media-specific hacks are not valid testing material, but then again composers will always try to stretch the limits of any given playback system. Any engineer trying to employ Ambisonic techniques will have to be aware of such corner cases.

Most DVD content showed a slight damping of the treble when played on Ambisonics. This effect usually caused a clear preference for discrete rendering, which disappeared after applying some corrective EQ.

One participant nicely summarized the overall consensus in remarking that in tape music, any form of playback is “interpretation”, and that he considered Ambisonic playback to be a different but valid interpretation except in special cases.

6. CONCLUSION

Informal listening tests have provided interesting insights into the potential and limitations of virtual speaker sources using Ambisonic panning. While acutely aware of the shortcomings, this author remains convinced that such systems are a very useful tool for electro-acoustic music facilities and can provide a very pleasing and successful concert experience. Since they can be built from commodity hardware and free software, the entrance barrier to a successful deployment is low, in theory. It is the author's hope that this paper contributes to lowering that barrier in actual practice.

7. REFERENCES

- [Adr10] Fons Adriaensen, the AMB plugin set, containing Ambisonic panners up to third order:
<http://kokkinizita.net/linuxaudio/downloads/index.html>, 2010
- [Adr10-2] Fons Adriaensen, AmbDec, a state-of-the-art Ambisonic decoder, l.c., 2010
- [Adr10-3] Fons Adriaensen, JAPA, the JACK audio perceptual analyser: l.c., 2010
- [Dav10] Paul Davis et al., Ardour digital audio workstation: <http://ardour.org>, 2010
- [Dav10-2] Paul Davis, Torben Hohn, et al., The JACK Audio Connection Kit: <http://jackaudio.org/>, 2010
- [Fol10] The CALF plugin collection, by Krzysztof Foltman and Thor Harald Johansen: <http://calf.sourceforge.net>, 2010
- [Ger74] Michael A. Gerzon: Surround Sound Psychoacoustics, in: Wireless World 12/1974:
http://www.audiosignal.co.uk/Resources/Surround_sound_psychoacoustics_A4.pdf, 1974
- [Lil10] Jonathan Moore Liles, Non-Mixer, an alternative JACK mixing/routing application:
<http://non-mixer.tuxfamily.org/>, 2010
- [Net08] Jörn Nettingsmeier, Ambisonics at Home:
http://stackingdwarves.net/public_stuff/linux_audio/tmt08/Ambisonic_Listening_Rig_with_Free_Software.pdf, 2008
- [Sbr09] Denis Sbragion, DRC, the digital room correction package:
<http://drc-firs.sourceforge.net>, 2010

8. ACKNOWLEDGEMENTS

The author would like to thank Martin Rumori and Judith Nordbrock at Kunsthochschule für Medien, Köln, and Prof. Thomas Neuhaus and Martin Preu at ICEM, Folkwang Hochschule Essen, for their kind support and expertise in conducting the listening tests.