3D SOUND FOR 3D GAMES - BEYOND 5.1

SIMON N GOODWIN

Central Technology Department, Codemasters Software Company, UK
simong@codemasters.com

7.1, HDMI 1.3 and 3D speaker configurations could potentially improve game audio beyond the scope of 5.1 cinema surround, but there are practical issues including ergonomics, compatibility and the lack of standard layouts. The paper explores possibilities and trade-offs and outlines a basis for standardisation. It seeks to justify the need for collaboration across the game and audio industry so that consumers benefit, rather than just get more confused. While there are still details that need to be agreed, the paper illustrates a set of approaches that give high quality 3D with equipment that is fast becoming the new consumer standard, while remaining closely compatible with existing CD and DVD mixes.

INTRODUCTION

The poster entitled 'How Players Listen' [1] set out the ways game audio reaches consumers, on PCs and consoles; all that was necessary before we can talk about improving what we've got. This paper goes beyond that towards things we don't do now but should, informed by research in the wider audio industry, the growing capability of game hardware and related systems, and the special needs of gamers.

The game industry is a bigger business worldwide than music or home cinema. Yet it has grown up in the shadow of music and video, repurposing hardware made for other purposes - like stereo music and home cinema, and largely independently of the audio industry and academia; this is after all the first AES conference specifically about gaming. Amplifiers are packed with features to tweak and remix music and video content, yet there's not one with a game button - even if the best thing that could do, most of the time, is turn off the chrome and confusing options, and let the game audio system speak for itself.

Games are arguably now the primary source of high-definition audiovisual content. Games regularly mix more sounds than big-budget films or albums; they're all interactive, and almost all make 3D graphics a key part of their appeal and fundamental to their gameplay. Yet there's no standard for 7.1 surround, let alone 3D loudspeaker audio - one attempt foundered a decade ago - and this conjunction of the AES and games industry gives us a chance to address that - and a creative and commercial opportunity for both.

1. HDMI OVERVIEW

This paper deals with HDMI, the High Definition Multimedia Interface, specifically for 7.1 speakers and HDMI 1.3; the lack of a horizontal standard; the 'hex plus' horizontal layout used in Codemasters games DiRT and GRID on PS3; the potential for 3D7.1, and some specific layouts we've tried.

HDMI supports eight uncompressed linear PCM audio channels with up to 24 bit resolution (144 dB) and sample rates to 192 KHz, but receivers and decoders made prior to the current HDMI 1.3 standard often treated audio as the poor relation, limiting actual capability to stereo or 5.1.

Sony's PlayStation 3 (with sales surpassing 10 million a year ago [2]) supports the full HDMI 1.3 standard, and Blu-ray disc players, the successors to DVD, make support for eight-channel LPCM mandatory - as well as the legacy DVD compressed 5.1 formats from Dolby and DTS.

Unfortunately the HDMI port on the most recent Xbox360 release, the Elite model, does not support LPCM beyond stereo, and delivers surround as encapsulated AC3, packed into the same space as the stereo output.

1.1 Analogue 8 channel audio

PCs have offered eight channel sound, typically configured as 7.1 (seven full-bandwidth channels and one 'LFE' dedicated to low frequency effects, derived from cinema convention), for more than a decade via multiple analogue line cables. The focus has been on filling the gaps in 5.1 audio, to the sides and rear, though Aureal pioneered an early 3D system.

These were analogue solutions, like the current X-Fi range from Creative Labs and its 7.1 predecessor the

Audigy, and separate from the video system, which developed from analogue VGA to digital DVI via separate cables, rather than a combined AV solution.

1.2 HDMI for PCs

But the PC is catching up with PS3 now; HDMI is becoming a PC standard as well as a video one; early PC HDMI solutions had limited audio, like the first-generation receivers, but Nvidia, ATI and Intel all offer uncompressed 8 channel HDMI audio output.

The ATI Radeon HD 4800 series of cards finally adds support for 8-channel LPCM output over HDMI, and NVIDIA's GeForce 8200 and Intel's G35 will output 8-channel LPCM over HDMI. [3]

2. CINEMA VERSUS GAMES

Some of the problems that limit cinema and music replay - the need to pack many customers between speakers, make a compromise mix to cover unknown listening conditions, catering for the worst-case at the expense of the best - do not apply to games.

Games are interactive and can be configurable or auto-configuring, since the player expects to make choices and explore and exploit the circular feedback between player and game. The player expects this to be fun.

Game audio can often justifiably be optimised for a single listener. It can adapt to known context - the dashboard settings in the console and game menus, but also the game mode - single or multi-player, network or split-screen, first or third person views.

But there are some problems too.

The first is that surround already had a chequered past. Layouts are ill-defined and widely ignored by consumers and content providers alike. Proposals go back to the 1930s - Alan Blumlein proposed surround sound with height by placing speakers above and below the screen in his 1934 patent [4]. It took 30 years for stereo to catch on, and as long again for 'home cinema', yet current surround systems are a temporary, fragmentary compromise, limited to two dimensions.

Michael Gerzon and kindred spirits researched flexible 2D and 3D surround systems in the seventies [5], and published intensively in AES and other papers, though as far as I'm aware Codemasters and Ubisoft are the only big game companies to follow this up, and we've still got catching up to do. Academics and artists have been playing with 3D sound for decades, but there's no breakthrough into the mass market, despite the extraordinary relevance of this to games.

The fact that patents were invoked in both stereo and Periphonics could explain the delay in uptake - as it did for high pressure steam in the 18th century - but that's no longer an excuse. There's no shortage of well-documented and freely available ways to render audio in two or three dimensions for any number of speakers, including VBAP and a slew of Ambisonics variants. I don't want to get into the holy wars of how we do it, or where the wires go - there are lots of solutions and it's a good thing for developers and equipment manufacturers to be able to experiment.

Surround in cinema is a nice but dispensible decoration. Side and rear channels add ambience and the odd spot effect, but those are optional extras - indeed mix guidelines say as much. But in games cues from beyond the screen can be a matter, almost literally, of life and death. If you don't know what's coming up behind you or to the sides, you can't block or avoid trouble, and only audio can provide those cues.

What I think we do need, analogous to the ITU 5.1 cinema standard [6] and informed by that, is a standard reference layout which better meets the needs of 3D games and in particular extends loudspeaker layout beyond a horizontal plane. This talk is a step towards that.

I'm not going to try to settle this in a talk at a conference. I can set out some of the parameters and priorities, as I see them - after discussion with academics, sound designers and game developers and looking forward to more - and ask if the AES consultation and consensus process can help. I and those I have consulted so far agree that having one standard, agreed and understandable, is more important than trying to get a 'perfect' standard, and we can progress only by agreeing a viable basic approach that can still be refined by the software and hardware development - rendering schemes, amp, speaker and stand design - oh yes and cabling solutions!

Perfection is unattainable - even if you're pragmatic, it's a moving target - but there is a strong case for 3D audio benefiting gamers - without making it compulsory - and once we've got it I suspect it'll be hard to live without, at least for audiophiles, and mass-market game applications will benefit niches including VR, teleconferencing, art and academia.

Innovation goes in several phases from dismissive incredulity to acceptance. Sooner or later we'll move through those and I'd rather it was sooner. Any 3D speaker arrangement for HDMI is likely to be transitional, as is ITU 5.1, but is a big step for gamers and audiophiles - indeed a step into a new dimension -

and the time is right to re-consider it now, with games a big business driving innovation, and HDMI providing convergence between high-def video and multi-channel coherent digital audio, and between PCs, dedicated game consoles and domestic media players.

Of course the AES has been here before. Alexander Von Humboldt reckoned there are three stages of scientific discovery: "first people deny it is true; then they deny it is important; finally they credit the wrong person." This has since been upgraded [7] to a five-step sequence:

1. People deny that the innovation is required.
2. People deny that the innovation is effective.
3. People deny that the innovation is important.
4. People deny that the innovation will justify the effort required to adopt it.
5. People accept and adopt the innovation, enjoy its benefits, attribute it to people other than the innovator, and deny the existence of stages 1 to 4.

Point 1 - games require 3D audio as much as they require better graphics, and arguably more so. Point 2 remains to be proven but those who have tried it are convinced it is at least possible. As regards point 3, it should not be compulsory but I think it's important, and the idea appeals to hardware manufacturers I've spoken with, as well as game enthusiasts and professionals.

Turning to point 4, the effort needs discussion - it's far less likely to be wasted if we work together, and it looks as if a lot of the work has been done already (in research, and equipment for horizontal surround). As regards point 5 I'm not looking for credit, just better audio, in games in particular, and the chance of fruitful collaboration with experts who feel the same way

3. PRACTICAL 3D AUDIO

At the moment the only way to get true 3D game audio is via headphone HRTFs, like Creative Labs CMSS-3D, or writing your own driver for a common API like OpenAL. Most game audio systems ignore the third dimension or try to map it into the other two, e.g. by distance effects. But all the information needed for 3D audio - positions, orientations, occlusion, reflecting surfaces - is available in every game engine, just waiting to be used.

Ideally we'll let consumers put speakers and their ears anywhere they like and work out how to render suitably. There are several ways to do this - beyond the current filtering and distance-adapting systems like Audessy, which ignore angles - but they involve directional or moving mics, and complex trade-offs between best and worst-case performance. That's a topic for another day, or another year. Right now let's find a way to prove the point for early adopters - luckily there are lot of them in gaming, and the industry structure encourages them among developers and consumers - and get into 3D.

Compatibility

The first and most practical and timely is to establish one standard layout for speakers not in a plane (i.e. for 3D rendering) which is still compatible with stereo and horizontal cinema 5.1 (for pragmatic reasons).

Agreeing upon a practical standard layout, with tolerances, is far more important than finding the 'best one' or specifying the rendering method. With eight standard uncompressed digital channels, HDMI is a wave we can ride, as AC3 was for 5.1.

Tetrahedrons

The simplest possible 3D layout is a tetrahedron - four speakers, three in a plane arranged as an equilateral triangle with one forming three more triangles beyond that plane. Figure 1 shows one possible implementation of this layout, with three speakers in front of the player and the apex of the tetrahedron behind.

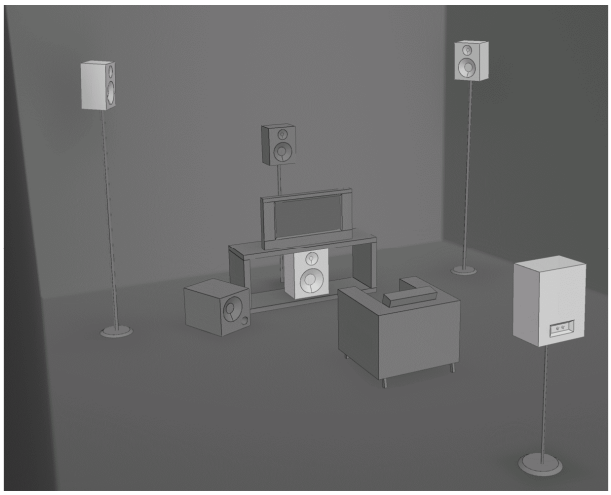

Figure 1

It's possible, but a bit less practical and far less compatible, to put one speaker in front and three behind – or three above and one below the listener, three below and one above, or various tilted variants.

I'm not advocating this layout, or any of the variants. It's theoretically possible to drive this with a 5.1 AC3 audio stream, leaving LFE and dialogue channels free for their conventional uses, but not compatible with existing 5.1 media.

While it could be compatible with existing hardware, it won't suit surround recordings expecting 5.1 speakers in the ITU layout - almost all existing media – without

substantial manipulation and compromises (e.g. combining surround signals for the single rear speaker) and it only suits stereo if the listener moves from the central sweet spot to sit on top of or preferably inside one of the speakers, so that two of the others form the required 60 degree spread. Then the putative 3D sweet spot will be half way down to the floor below that…

If you put the plane at the top you need a speaker right below you, which isn't very convenient - three speakers on the floor are at least easy to line up but compromised by occlusion and pressure zone effects, and hanging the fourth above you isn't easy either.

Most importantly, Gerzon was unable to make the tetrahedron work as a convincing way to render 3D audio, and analysed why with characteristic wit and rigour [8]. After dismissing a 20 speaker dodecahedronal arrangement with the sentence, "The reason for the impracticality of this is self-evident to those whose living rooms are not anechoic chambers with wire mesh floors" he considered the tetrahedron, as the simplest regular polyhedron potentially capable of treating all directions as equivalent, but concluded that 'defects are inherent in the tetrahedral speaker layout' because, 'the Makita and energy vector locations do not coincide… sounds at high frequencies are very much drawn towards the four loudspeakers … other speaker layouts must be used."

We can't in practice roll in the other one and a bit speaker feeds of movie surround to help out, because 5.1 standards and 7.1 conventions reserve two channels for special cinema-based purposes. The optional LFE channel is severely bandwidth limited by DTS and Dolby Digital compression (but not HDMI) and intended only for low frequencies - in practice its use is bound up with bass management and the drive to make other speakers cheaper by diverting massive bass effort to that speaker. The fifth channel is notionally full-bandwidth - though often implemented with a special speaker - and dedicated to speech or dialogue.

Games may use this to fill out panning across the front but ours - like many - follow the cinema convention and use it only for sounds not placed in the world, like speech and (in the case of games) menu interface beeps, confirmation and cancellation signals. Problems may arise if that speaker is used to render the sound of objects in the game world as it may be a different type, with characteristics optimized for speech, and the location - above or below the screen, and ideally inside it - is unpredictable.

So we're back with four speakers and a couple of odd ones in 5.1, typically compressed and multiplexed by systems that rely on psychoacoustic masking and don't preserve phase. There may yet be a way to offer convincing 3D over AC3, but I've yet to hear, or imagine, it, and decades of research (since Pierre Schaeffer in 1952, cited in [8]) suggest otherwise.

But AC3 is old hat, and in HDMI 7.1 we're left with six plus the odd couple. In theory HDMI could offer full-bandwidth over eight channels - 8.0 rather than 7.1, ideal for a cube - but current amplifiers treat the .1 channel specially, with different circuits and filtering, and I'm strongly advised by Arcam, Denon and Onkyo developers not to try 8.0. Though it's feasible for amplifiers to be redesigned without changing HDMI, I want to work with what exists, not wait for new hardware, and a cube is problematic for compatibility with the thousands of DVDs and games that use the odd couple of channels specially.

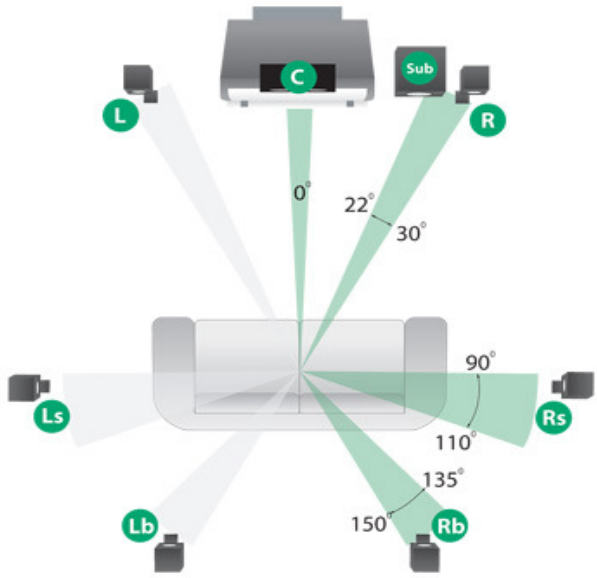

Figure 2

Horizontal 7.1

Back in 7.1, there's still six channels available for directional information. Amazingly there's no equivalent of the ITU standard for 7.1, but Dolby, Creative Labs and THX have de-facto standards and implementations which have reached a sort of consensus for horizontal 7.1. The big gaps in 5.1 are at the sides and rear, so the approach is to fill out the sides with the two extra speakers and move the back pair further round to cover the rear better. The front spread remains +/- 30 degrees, for Blumlein stereo compatibility, plus front centre in there somewhere. The other angles are a compromise between 5.1 and a regular hexagon, with Dolby recommending a range that encompasses both (Figure 2, [9]) and THX opting for something similar but with control over the exact angles and distances (Figure 3,

showing the THX Setup Console bundled with X-Fi PC audio cards).

Figure 3

These angles work well. Our games DiRT and GRID use a compliant regular hexagon, via hybrid third order Ambisonics on PS3 [10] and Creative Labs own proprietary panning on PC. This gives a much wider sweet zone and more freedom in speaker and listener placement than 5.1, but still only two dimensions.

We found our vector-optimised hybrid higher-order Ambisonic panner runs faster than the default 5.0 VBAP implementation in Sony's SDK. Our approach is similar to that outlined in Dave Malham at the AES UK 11th Conference [13] but using third order rather than second order for the horizontal components. The vector architecture of the PS3 means we get the extra two channels with no extra processing overhead, and improves panning of 5.1 output to ITU angles, but the first-order elevation calculations are of little benefit to the player if the speakers are all in one plane.

In theory we could implement a custom configuration screen to allow users to specify the azimuth and elevation of each speaker – though this leaves the vital question of the 'best' default unresolved. A horizontal equivalent control appears in a few games, such as the Microsoft-published Project Gotham 3 and 4 for Xbox360 (for fewer speakers), and in the PC setup tool in Figure 3, but so far that fully-manual approach has been avoided on the grounds of complexity and the risk of user error making things worse rather than better. To simplify user setup we'd all benefit from agreeing one practical default 3D layout.

Aureal Vortex cards pioneered hardware HRTFs and 7.1 sound in the 1990s, but their 3D arrangement - adding two low front speakers to give a W layout – did not take off. This kept the wiring simple but arguably accurate spatialisation is least relevant in front of the player, where the visuals take precedence, and was deficient for sounds beyond the screen, at the sides and rear, where audio could add most to 3D games.

So the first thing I tried to get a game into 3D - and implemented in Colin McRae DiRT, though as a hidden option inaccessible to consumers - was dropping the side channels out of the plane of the 5.1 set. This gave some 3D cues but was biased towards the sides, where it's not as useful as vertical information to front and rear in most games - as those tend to be where you're going and coming from - and a dip in level it was hard to code away as sound sources moved past on either side.

Gerzon's papers include promising research into true 3D sound rendered via a regular octahedron [8]. On a trip to York University MRC to attend the SpACE-Net conference in January 2008 I spotted six speakers arranged in two triangular planes - pointing in opposite directions - set up for an experimental 3D audio art installation [11], and discussed that with Tony Myatt and Dave Malham. They weren't aiming for technical fidelity, but found the arrangement gave good cues all round, up and down.

An earlier set-up at York was less satisfactory – the octahedron of Figure 6 in [8], with left and right side speakers and high and low pairs in front of the listener and behind – was found to be unstable in practice, working for a listener in the central sweet spot but not well otherwise. It seems that the most effective orientations of the octahedron have speakers left and right of the listener both in front and behind, and we speculate that this is because humans have left and right ears, rather than top and bottom ones, and thus greater sensitivity to spatialisation in the horizontal plane. This also gives better chance of compatibility with existing stereo.

I suspect - despite the ITU/game quad angle discrepancy – that compatibility is so vital there's little scope for a proposal that does not do a good job of rendering existing mono, 1D (stereo) and 2D (ITU 5.1) media, and ITU doesn't look much like an octahedron, from any angle. I then tested some speaker layouts at work, and started to try to reconcile the needs of ITU and a possible 3D HDMI layout, which I propose we call 3D7.1. If CDTV and I3DL2 were proper names, that's as good, I reckon, and doesn't presume any one topology or brand. My experiments are promising and could inform others to make further improvements, so here's what works for me, and why.

At first I went for a tilted arrangement, letting the rear/surround speakers drop with one placed high - rather higher than the front pair - and one at the front, between and below the usual stereo pair, now raised.

This pans game audio as well as quad, either with VBAP panning or hybrid Ambisonic code, but in three dimensions as long as the speakers are not occluded and the listener is centrally placed, as shown in Figure 4.

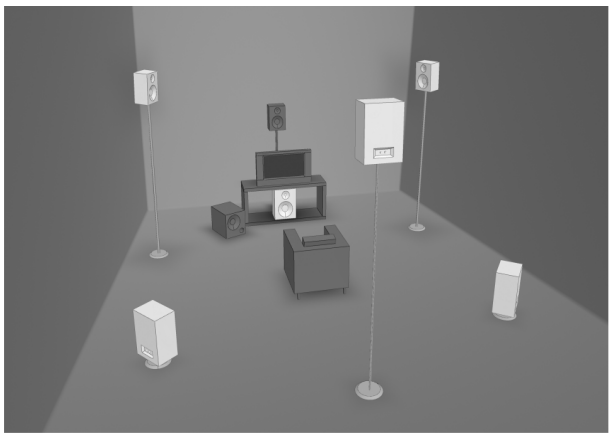

Figure 4

Table 1 shows the azimuth and elevation angles in degrees for the six main speakers and Cartesian co-ordinates on a unit sphere where X is front/back (+ forwards), Y is left/right (+ left), Z is up/down (+ up) and azimuth is anti-clockwise from a view facing the font speaker and elevation is positive for angles above the listener (consistent with previous 3D audio papers).

Speaker position	X	Y	Z	Azi	Elev
Front Left High	**+0.408**	**+0.707**	**+0.577**	**45**	**+35**
Front Right High	**+0.408**	**- 0.707**	**+0.577**	**315**	**+35**
Rear Left Low	**-0.408**	**+0.707**	**- 0.577**	**135**	**- 35**
Rear Right Low	**-0.408**	**- 0.707**	**- 0.577**	**225**	**- 35**
Rear Centre High	**-0.817**	**0**	**+0.577**	**180**	**+35**
Front Centre Low	**+0.817**	**0**	**- 0.577**	**0**	**- 35**

Table 1

This scheme gives acceptable results for most cinema surround, as the rear channels are typically used for ambience, especially if the listener is sitting further back - bringing the rear speaker angles closer to ITU range (+/-100..120 degrees) and the front pair from the constant-power de-facto standard of +/-45 most common in games to the Blumlein-favoured +/-30 that avoids a hole appearing in the centre of a stereo image.

Indeed there is a stereo-friendly sweet spot at the back of this rig, a good match for Dolby home cinema guidelines just in front of it, and an ideal location for gamers further forward, and it's natural, furniture permitting, to move between those as you listen if you value exact balance. Figure 5 shows the layout suitable for listening to pre-rendered cinema soundtracks – the only difference is the preferred listener position.

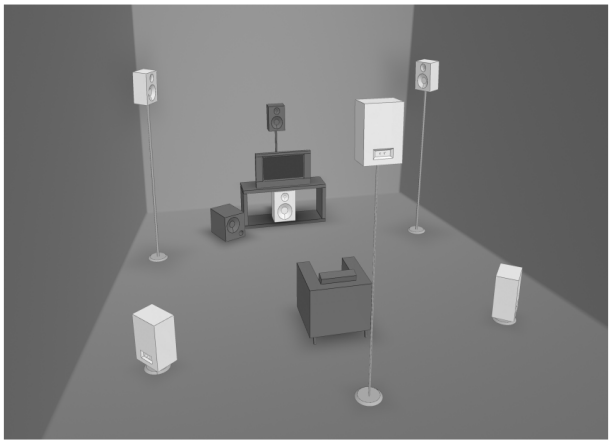

Figure 5

Empirically it seems reasonable to say that players lean forward or sit closer to the screen than film watchers, even without quantitative evidence to 'prove' this. In any case this concept of a sweet range, optimised for particular tasks, is worth developing, because it allows more creative compromises and flexibility than conceptually nailing the listener to one spot, and the listener can naturally and interactively prioritise each to meet their changing needs. It's a lot easier to move your head, or a chair, than eight speakers and a big television.

When I discussed 3D7.1 with Mark Yonge, at the April 2008 Cambridge UK conference of the AES, he identified another variable within the constraints of periphonic research and compatibility with existing material. Mark proposed rotating the polyhedron so that all the rear speakers are moved higher, reducing the risk of them being occluded and potentially bringing the front speakers down a bit, closer to the preferred height for cinema and stereo, and the front and rear triangles towards planes perpendicular to the floor, which is convenient in rooms with walls arranged the same way. Figure 6 shows this layout.

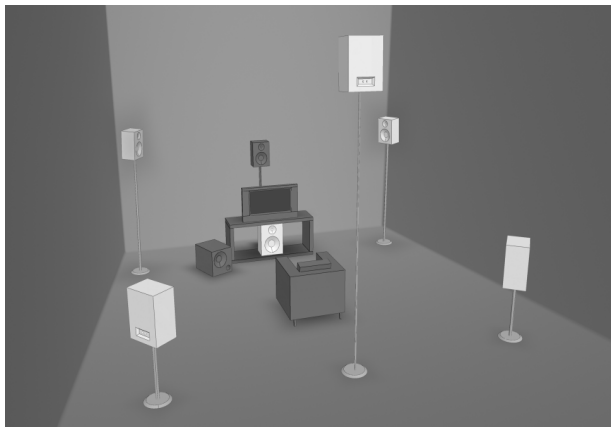

Figure 6

It's still not ITU compliant - surround speakers are too low - but is close; a substantial improvement, with the tilt just 24 degrees (ITU specs permit surround speakers 15 degrees down) and in the direction that a relaxed listener is likely to flop in. The ITU recommendations seem to assume cinema-goers sit bolt upright…

The banked seating in cinemas implies either a forward tilt down for speakers and listeners or means that only a few rows of listeners experience ITU-compliant elevations. Such nuances seem best left to THX. Table 2 shows the co-ordinates and angles, according to the same convention as Table 1.

Speaker position	X	Y	Z	Azi	Elev
Front Left High	+0.577	+0.707	+0.408	45	+24
Front Right High	+0.577	- 0.707	+0.408	315	+24
Rear Left Low	- 0.577	+0.707	- 0.408	135	- 24
Rear Right Low	- 0.577	- 0.707	- 0.408	225	- 24
Rear Centre High	- 0.577	0	+0.817	180	+55
Front Centre Low	+0.577	0	- 0.817	0	- 55

Table 2

The octahedron has been rotated 90 degrees (with speakers re-labelled) so component triangles are parallel with front and rear walls rather than floor and ceiling. I contend that this tilt does not significantly detract from the listener's experience of audio mixed in 5.1 for cinema consumption, and seems at least as likely to be beneficial as detrimental. More tilt tends to bring the top back and bottom front centre speakers in towards the listener, which could cause ergonomic problems…

The ideal angle is a matter for discussion but so far I favour wall-parallel front and rear triangles, for pragmatic reasons. Figure 7 shows the room layout which works best when listening to pre-rendered 5.1 with front and rear triangles of the octahedron in two vertical planes. It is identical to figure 6 except that the listener's seat is placed further away from the screen.

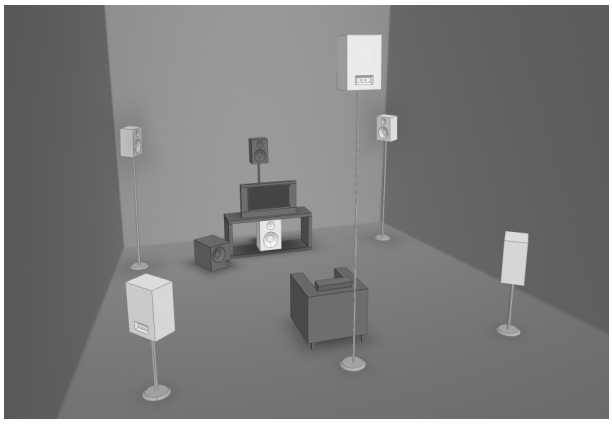

Figure 7

Even if 3D7.1 simply opts for a regular octahedron, leaving two channels for front centre and LFE, we need to decide and justify its orientation, angles, tolerances and where the listener goes.

CONCLUSION

Speakers arranged in a conveniently orientated Octahedron, plus LFE and Front Centre units, deliver good compatibility with existing stereo and 5.1 media while offering the potential for true periphony on HDMI 1.3 compatible hardware.

This arrangement has been implemented by Codemasters on Sony PS3 via HDMI, using the same code as in the million-selling DiRT and GRID games – with reconfigured speaker angles – and found to work well. It is an option on the new OpenAL driver from Blue Ripple Sound [12], making it available on the PC versions of those games (without recompiling, as the driver is dynamically bound) and similarly on very successful games such as Unreal Tournament, Gears of War, Battlefield 2 and 2142 and many others.

HDMI interfaces, OpenAL drivers, and Codemasters games can all help to make 3D7.1 a reality, but the chance is best if there is wide discussion and collaboration involving academics and others in the wider audio industry. The AES seems well placed to facilitate this, and ensure the benefit extends beyond gamers, big though that market now is.

In discussions at the 2007 Game Developers Conference in San Francisco, Creative Labs and Sony Audio managers George Thorn and Jason Page indicated that they would support a 3D7.1 audio standard by providing a compatible option for users, if consensus on a suitable layout could be reached. The idea was also positively received at the meeting of the Interactive Audio Special Interest Group (IA-SIG) at the same show. But nothing will happen without agreement and AES members are well aware that multiple 'standards' must be avoided.

The case for implementing 3D7.1 with an octahedron plus front centre and LFE channels seems compelling. Another useful empirical observation is that game players sit further forward so the angles that are right for ITU, sitting back, morph to symmetrical periphony as players lean forward and move towards the centre.

Some technical questions remain, besides the obvious political ones. What front/rear tilt is acceptable? Can the rear top speaker (and front low one) compensate for ITU 5.1 tilt, perhaps in a manner similar to the 6.1 approach of DTS ES Matrix? Perhaps the existing front centre, and extra front top (or bottom) speaker in an

octahedron, can help to solve the 'speakers in the screen' problem, creating an unoccluded centred image?

Is there scope for amplifiers exploiting an agreed regular arrangement of speakers in an octahedron to synthesise feeds for other arrangements, perhaps by the techniques outlined by Laborie, Bruno and Montoya [14]? This could reduce the need for consumers to fit their furniture around one speaker layout, without the conversion requirements of direct B-Format or higher order component output. Output for an orthogonal array of speakers can be reconstituted via spherical harmonics, for adaptive repositioning, relatively simply.

While Ambisonic techniques have been cited here, there's no reason why VBAP, Ambiphonics, HRTFs or other techniques could not be used to derive speaker feeds for a 3D7.1 rig. It's important that this discussion of a layout should not get hung up on rendering schemes, where there's scope for experimentation and innovation.

Is there production-engineering benefit in driving the array from four amplifiers, rather than the six currently required, as Gerzon suggested [8]? How best can users arrange the furniture, including speaker stands, ceiling or wall mounting for high speakers, and where do all the wires go?

Making the best of the opportunity to implement 3D speaker audio, spurred by the potential of HDMI and demands of games, means working towards a consensus with a multi-background team while leaving them all enough scope to make their version unique for market reasons while close enough to the 'standard' for a consistent consumer benefit. Can you help?

REFERENCES

[1] Goodwin, Simon N., "How Players Listen", AES 35th Conference, February 2009

[2] PlayStation 3 puts HDMI 1.3 in 10 million homes: http://www.gamesindustry.biz/articles/playstation-3-sales-hit-105-million

[3] ATI HD 4800 and similar HDMI 1.3 PC AV cards: http://www.anandtech.com/video/showdoc.aspx?i=3338&p=2

[4] Blumlein, Alan D, "Improvements in and relating to Sound-transmission, Sound-recording and Sound-reproducing Systems", JAES Volume 6 Issue 2 pp. 97; April 1958

[5] Gerzon, Michael A, "Periphony: With-Height Sound Reproduction", JAES Volume 21, pp 2-10; January 1973

[6] ITU-R BS 775-1, cited in "Multichannel surround sound systems and operations", AES Technical Council document AESTD1001.1.01-10, pp. 4; 2001; online at http://www.aes.org/technical/documents/AESTD1001.pdf (checked 2008-11-19)

[7] Fish, Syd (pseudonym), "Five stages of innovation" http://sydfish.wordpress.com/2008/03/17/the-five-stages-of-innovation (checked 2008-11-19)

[8] Gerzon, Michael A, "Practical Periphony: The Reproduction of full-sphere sound", AES 65th Convention, February 1980

[9] Dolby 7.1 speaker setup guide, online version at: http://www.dolby.com/consumer/home_entertainment/roomlayout.html (checked 2008-11-18)

[10] Delefie, Etienne, Ambisonia web site, http://ambisonicbootlegs.wordpress.com/2007/08/30/interview-with-simon-goodwin-of-codemasters-on-the-ps3-game-dirt-and-ambisonics/ (checked 2008-11-19)

[11] Thyssen-Bornemisza Art Contemporary surround sound installation, Vienna, Austria, http://www.tba21.org/program/current (checked 2008-11-18)

[12] Furse W. E., "Building an OpenAL implementation using Ambisonics", AES 35th Conference, 2009

[13] Malham, David, "Ambisonics in the New Media-'The Technology Comes of Age'", AES UK 11th Conference, pp. 99; March 1996

[14] Laborie, Bruno, Montoya, "Reproducing multichannel sound on any speaker layout", AES 118th Convention, May 2005

Originally published on behalf of Codemasters Software Ltd and Simon N Goodwin by the Audio Engineering Society. Last revised 13th September 2013. Updated on 25th August 2009 for general distribution under the Creative Commons Attribution-Share Alike 2.0 UK: England & Wales licence: