

IA-SIG Newsletter
"The Interactive Audio Journal"
www.iasig.org

Vol. 1 No. 1, March 15th 1999
Editor: Alexander Brandon

In This Issue

Section I: From the Chairman — IA-SIG Chairman, Mark Miller, discusses the IA-SIG's 1999 goals

Section II: Official Announcements — All official announcements regarding IA-SIG members, activities, and special events. In this Issues, the IA-SIG at GDC and IMX.

Section III: Working Group Reports — Status Reports from the 3D Working Group, the Interactive Composition Working Group, The Audio Advisory working Group, The Platform Development Working Group and the Multi Format working Group.

Section IV: Features — This section contains features and columns on varied topics of IA-SIG interest. This issue features our first Interactive Composition Column, "A Direct Music Primer" by Alexander Brandon and Mark Miller

Section V: Industry Corner — Tools and technology development are the driving forces behind the realization of musical creativity. This section features reviews and directories of useful tools, technologies, and applications. This issue we have an "Engine Roundup" directory of Interactive Music Engines by Alexander Brandon

Section VI: Developers Corner — At last, a place where the developers can speak their minds. The interactive audio industry is bursting with men and women ready to both complain and cheer about various aspects of their work. Here is an outlet for them. This issue features comments from Marty O'Donnell, Chris Grigg, David Javelosa and Spencer Critchley

If you are interested in contributing to "The Interactive Audio Journal" please contact Mark Miller (mark@harmonixmusic.com)

Please Join the IA-SIG!

The Interactive Audio Special Interest Group (IA-SIG) exists to allow developers of audio software, hardware, and content to freely exchange ideas about "interactive audio". The goal of the group is to improve the performance of interactive applications by influencing hardware and software design, as well as leveraging the combined skills of the audio community to make better tools.

The IA-SIG has been influential in the development of audio standards, features, and APIs for Microsoft Windows and other platforms, and has helped numerous hardware companies define their directions for the future. Anyone with a commercial interest in multimedia audio is encouraged to become a member of the IA-SIG and participate in IA-SIG discussions. (Visit <http://www.iasig.org> for membership details).

IA-SIG Steering Committee

Chairman: Mark Miller (mark@harmonixmusic.com)

Steering Committee Members: Rob Hubbard (EA), Monty Schmidt (Sonic Foundry), Danny Petkevich (Staccato systems), Brian Schmidt (Microsoft), Alexander Brandon (Straylight Productions), Tom White (MIDI Manufacturers Association).

IASIG Advisory Board

Thomas Dolby Robertson (Headspace), David Mash (Berklee), Craig Anderton (EQ), Gordon Currie (Portal Productions), Dale Gulick (AMD), Rudy Helm (Helm Productions)

IA-SIG Glossary of Terms

- "DLS" — DownLoadable Sounds, this is a standard created by the IA-SIG to allow users of computer audio cards and software synthesizers to create their own instruments and sounds rather than be limited to a specific sound set.
- "WG" — Working Group, the working end of IA-SIG, these groups represent various causes in the interactive audio industry. They work towards goals that might not otherwise be possible in any other body.
- "GDC" — Game Developers Conference, held once each year, this conference brings together some of the best in the game industry for both computers and game console systems to discuss current issues, teach classes, and explore new opportunities for gaming.

Section I: From the Chairman

"IA-SIG in 1999"

With the GDC right around the corner, I thought that I would drop you all a note about the IASIG's plans for '99. As I see it, this is a most critical year for us. While we have had some great recent successes with DLS and 3D Audio, the pace of change in our industry is accelerating. Our standard work is even more crucial and important now than at any time in the past. With the help of the Steering Committee and the new Advisory Board, I have set the following strategic priorities for our organization in the coming year.

Where we have succeeded in the past, we will continue. Where we continue, we will improve.

The IASIG has been very successful in bringing the industry together and solving some of our most vexing problems:

- DLS I is finally entering the main stream
- DLS II will unify Sound Fonts and DLS and gain even wider adoption through its inclusion in the forthcoming MPEG 4 standard
- Hardware accelerated interactive 3D audio is working well in DirectX and is being used in almost every cutting edge game published
- We plan to announce the Interactive 3D Audio Level 2 specification at GDC.

While we have occasionally taken longer than we would have liked, none can doubt the impact of our work. Into the new year, we will continue to form relevant Working Groups and create meaningful and useful Standards and Practices documents. In addition, we will endeavor to improve our operations in the following ways:

1) Get the word out!

The IASIG is in the information business. The standards and practices that we create through the WG process are our products. The people who create audio content or audio technology for interactive platforms are our customers. In the coming year, our main goal will be to improve the efficiency and scale of the distribution of our products to our customers. To achieve this, we will undertake a major new initiative.

Beginning in March of '99, the IASIG will publish (via email) a quarterly newsletter containing the news and information that is of vital importance to our industry. Each issue, each Working Group will publish their quarterly results and findings along with as much people, product, and technology news as we can gather. My hope is that this publication will quickly become the industry's main source of information on Interactive Audio.

Alexander Brandon will be heading up this effort. Anyone who is interested in becoming involved or who has news that is 'fit to print' should contact Alex <alex@mail.portup.com> or myself directly.

2) Target the adoption cycle of software developers.

New audio technology and features are great. The innovations and evolutions of each successive generation of hardware and software audio technology is the life's blood of our business. However, if these new features are not supported by software developers, chances are, no end users will ever enjoy the benefit. If the consumer sees no tangible reward for upgrading the platform, our market will wither and all will suffer. Yet, if the software development community adopts each new level of technology uniformly as it emerges into the marketplace, the upgrade proposition will be a very attractive one to the end user and all will prosper.

It is our mission as the IASIG to come to agreement on how to move the consumer and developer forward and then package that information in a single, simple 'upgrade' message. Using our newsletter, our other journalistic efforts, conference venues like the GDC and our ongoing work with the providers of audio SDKs, we will dramatically increase our efforts to educate the software development community and consumer and facilitate their adoption of new audio technology.

3) Strategically take the long view where possible.

Collectively, we represent many of the best minds in the audio industry. While it is often difficult to spend our time thinking beyond the present fiscal quarter, it is perhaps the most valuable contribution we can make. For example, three years ago, the Interactive Composition Working Group took on just such a challenge. We created a lexicon and began to explore the outer limits of Interactive Composition well before it had much of any main stream importance. As a result, we were ready when DirectMusic was being born and many of our terms and ideas were subsequently adopted (for the benefit of all involved.)

In this year, I want to encourage those of you who have the vision (or the luxury) of investing in the exploration of the next two to four years, to initiate more forward thinking discussions and Working Groups. If we can execute on these points, I believe that we will rise to a new level and come fully into our own as an organization.

I am grateful to all of you for all of the hard work that you have done to date and look forward to a great new year.

Yours,

Mark Miller
Chairman

Section II: Official Announcements

"I3DL2"

For Release March 18th 1999

"Industry leading companies pledge support for enhancements to Interactive 3D Audio."

This week, Creative Labs, Aureal Semiconductor, QSound Labs, Sensura, Conexant, VLSI Technologies, Yamaha, DaimondWare, and others pledged support for the Interactive Audio Special Interest Group's (IA-SIG) open guidelines for Interactive 3D Audio Level 2 (I3DL2) enhancements. The guidelines define highly realistic, interactive simulations of spatial acoustics and real time audio environment modeling. I3DL2 enhancements will allow for much more realistic and immersive 3D entertainment experiences for the PC and other digital entertainment platforms.

Mark Miller, the IA-SIG Chairman had the following to say: "We are very proud to announce support for these new enhancements. It is through exactly this type of industry cooperation that great progress can be made. The I3DL2 enhancements will provide tremendous benefit to software developers, hardware manufacturers and most of all consumers by increasing the realism, consistency and quality of Interactive 3D Audio experiences across all platform."

The enhancement specification was drafted primarily by Jean Marc Jot of Creative Labs with critical support and input from the entire 3D Audio Working Group (3DWG), and is based on Creative's Environmental Audio Extensions (EAX) to Microsoft's DirectSound3D. Later this year, the 3DWG will finalize the document and publish a platform-independent version of the I3DL2 guidelines. The companies contributing to the working group process (including Aureal Semiconductor, Creative Labs, QSound Labs, Sensura, VLSI Technologies, Conexant, Microsoft, Spatializer, Yamaha and others) are all members of the IA-SIG 3DWG.

Last year, the IA-SIG published a set of guidelines known as 'The IASIG Interactive 3D Audio Rendering and Evaluation Guidelines, Level 1' (I3DL1). This document, based largely on Aureal Semiconductor's A3D white paper, includes a primer on 3D audio, guidelines for accurate auditioning and evaluation of 3D audio performance, and a specification which describes the minimal requirements for I3DL1. The specification is accompanied by Minerva, a 3D audio performance testing program written by Aureal and donated to the IA-SIG for industry usage. Industry support for the I3DL1 guidelines has resulted in a dramatic increase in the use of positional 3D audio in a wide variety of games.

For more information, please visit www.iasig.org or contact Mark Miller at mark@harmonixmusic.com.

"IA-SIG and IMX join forces"

(IMX) www.imusicxpo.com is proud to welcome the MIDI Manufacturers Association (MMA) and IASIG as its newest Association Sponsors! The MMA/IASIG, its member companies and IMX share an important common goal; to unite, explore and demonstrate how technology is influencing and empowering end-users and creators of music.

James Grunke, Chairman of the MMA recently commented at an association meeting, "We feel the IMX show fills a void for the MI industry by providing a world class showcase for music technology vendors to reach out directly to their customer base, musicians."

Commenting on the debut of the IMX event, ISG President Paul Sitar said: "Without a doubt, the internet and new technologies are the most important influences music has experienced in recent history. There are other forums but none with the purpose for and a clear focus on musicians. It all starts with the creators/musicians. Without them, there is no music industry. IMX's mission is to help understand and explore interactive music."

Grunke further states "We are excited to be a sponsor of the IMX show, the MMA/IASIG represents the cutting edge of music technology companies and we fully expect this event to become an important part of our members annual show schedule."

Complete, up-to-date information can be found at <http://www.imusicxpo.com>.

"The Audio track at GDC 1999, first reports are in..."

Audio is no longer taking a backseat to graphics in the game industry, and this year's GDC turned out the largest representation of audio vendors and developers yet. Attendance was also made available through a discounted "Audio only pass" that allowed attendees to come to just the interesting (audio related) events at the show. The AudioPass was a huge success as was the conference audio track over all. Near as I can tell, we have around 250 people coming to the Audio sessions. In fact, attendance was so strong (and unexpectedly so) that our main session room had to be reconfigured to make room for all of the people who wanted to get in. They actually had to make another, non-audio conference room smaller to accommodate us. This is all the more amazing based upon the fact that there were only 60 people registered for audio in 1998. In 1998, most of the session had only 20 people or so and it was really kind of depressing for both the speakers and the attendees.

In fact, I actually arrived at this year's show a day early to scope things out. When I saw that the audio room seated over 100 people, I got very nervous. 20 people in a 100 seat room is even more depressing than 20 people in a 50 seat room. I almost went and asked the conference organizers to move us to a smaller room, but then I caught myself. We had made many changes to the track for this year (which I will detail below) and it would have been cowardly to not give the expanded venue a go.

As it turns out when I arrived at the first session (a notoriously poorly attended slot for audio people), I was quite surprised to find the door closed. My first thought was that no one had come so the session had been cancelled. My second thought was that my watch was running slow and that I had missed most of it. The truth of the matter was that the room was full. Full to standing room capacity and the doors had been shut for reasons of fire safety.

I did finally get in to introduce the track and the speakers, but for the entire rest of the day, the door has to be propped open and an extra speaker positioned out in the hall so that people could stand outside and listen. That night they reconfigured the room. The next day we continued to fill the new 180 seat space and a 40 seat roundtable room, every session, for the remainder of the show. Simply amazing!

Last year, my first year on the GDC Advisory board, I had done a lot of experimenting with the audio content. As it turned out, many of these experiments failed. As a result, the track was poorly attended and fairly unimpressive. Most sessions had no more than 20 people attending. Basically, it was bad but it taught me some valuable lessons.

Most importantly, I took away the following ideas: The people who create audio and audio technology for games comprise a community. Those inside the community are generally highly skilled and well informed about the state of the art. 'Audio 10' classes for these people are a grand waste of time and money. Instead, very dense, up to date and specific content is required make their experience of the conference compelling. Those outside of the community, while often well meaning, really don't care to know about that daunting level of detail. All they (mostly young producers) really want to know about how to deal with and manage those unruly audio people.

The other big idea came from attending George Sanger's BBQ (in his words 'an intense Texas style think tank'). This idea came from the realization that most of the expertise in any given session was in the audience, not behind the podium. Getting the attendees involved in the knowledge sharing process is, therefore, an important key to a successful conference.

I tried to build the conference sessions around these two ideas. The result took the form of a single, linear track. The sessions would be divided into three types; Master Classes in specific skill areas, technical sessions on topics of current interest, and roundtable discussions on related topics. The roundtables were scheduled to follow related sessions in order to give every one a chance to talk openly about the information they had just consumed. Session abstracts can be found at: <http://www.gdconf.com/cgi-bin/1999/main.pl?x-x=51497> and a detailed description of what went on will follow in the next issue.

All of the IA-SIG meetings at the show (detailed below) were well attended, but our big news came on the front page of the show daily on Thursday. The 3DWG had produced substantive agreement on the IA-SIG Interactive 3D Audio Rendering Guidelines, Level 2. Simultaneous press releases were issued by Creative Labs, Aureal Semiconductor, Sensaura, Q-Sound and the IA-SIG touting this historical event. You most likely know about this already as the SIG release is printed above, at the beginning of this section.

Overall, it was the most exciting and fulfilling event that I have ever had the privilege of being professionally involved with. I will follow this up in the next issue with a more detailed report but I had to take this opportunity to share my excitement about it with you in this, our first issue.

Lastly, let me once again thank Microsoft's DirectMusic team for sponsoring our general meeting and Yamaha for supplying an outstanding 5.1 speaker system.

Yours,

Mark Miller

GDC Meeting abstracts:

IA-SIG General Meeting

The IA-SIG General Meeting will be held at the 1999 Game Developers Conference in San Jose, CA on Wednesday, the 17th of March from 6:00PM until 7:30PM in room B 1/4 in the Convention Center.. Courtesy of Microsoft's Direct Music team, our most generous sponsor, there will be light refreshments and cocktails after the meeting.

Meeting Agenda:

- State of the SIG address (outlining our goals for the coming year)
- Introduction of the Steering Committee and Advisory Board Members
- News of the recent efforts to promote a Grammy award for Video Game music from Chance Thomas
- Message from Joe Paulino, Vice President of the San Francisco Chapter of Screen Actors Guild
- Progress report from the 3D Working Group
- Reports from the Audio Advisory Working Group, the Interactive Composition Working Group, and the newly formed Multi-Format Working Group
- Presentation of inaugural IASIG news letter by Alex Brandon, editor.
- Open discussion on topics of your choice (For example, 'the major hindrances to audio developers')
- Message from the meeting sponsor, the DirectMusic Team

3DWG Meeting

The 3D Working Group will hold a meeting on Thursday, the 18th from 2:00PM to 3:00PM also in room B1/4. The main topic will be the 'IA-SIG Interactive 3D Audio Rendering Guidelines Level 2'. Existing members and all interested parties are encouraged to attend.

AAWG Meeting

The Audio Advisory Working Group will hold an open meeting on Friday, the 19th from 2:00PM to 3:00PM in room B1/4. The main topic will be their plans for bringing greater awareness of the value of high quality audio on the PC desktop to industry decision makers. Existing members and all interested parties are encouraged to attend.

Special Session

Also, please plan to attend 'Interactive Audio: What the Future Holds'. This crystal ball vision of the future of interactive audio will be presented by Rudy Helm and Rob Hubbard, the previous Co-Chairmen of the IA-SIG, and will feature some very interesting guest speakers. This session will be held from 5:00PM to 6:00PM on Wednesday, the 17th in room B1/4. (just before the General Meeting).

The IA-SIG has also arranged many round-table discussions and panels related to issues of interest to audio developers which are part of the GDC conference program.

Section III: Working Group Reports

=====

The Working Group is the main functional aspect of the IA-SIG. Working Groups generally form around issues of current concern for the industry. Once formed, they meet either in person or via the Internet and develop standards and recommended practices document. These documents represent industry consensus and are published and made available to all interested parties. This is where the Working Groups (WGs) report their quarterly progress. As this is a first issue, an overview of each group will be presented to introduce them. In future, full descriptions of WG functions will be available on the website.

3D AUDIO WORKING GROUP (3DWG)

Chairman: Conrad Maxwell, Conexant <conrad.maxwell@conexant.com>

Overview

The group is focused on creating 3D Audio rendering standards to define more realistic audio environments. This effort has led to extensions to the Microsoft DirectSound 3.0 API to enable hardware acceleration, and to the publication of the IASIG Interactive 3D Audio Rendering and Evaluation Guidelines Level 1, describing "minimal acceptable" 3D audio features for all platforms. The group is now discussing standard enhancements to current 3D audio technology, such as Reverb parameters, object reflections and occlusions, and more, for release as Level 2 Rendering Guidelines.

Current Status

The 3D Working Group is nearing its second major milestone-- the ratification of Interactive 3D audio Rendering Guidelines Level 2.0. Level 2 attacks the problem of enhancing the 3D audio experience with environmental reverberation, occlusion and obstruction modeling. The Level 2 extensions create a common property set and behavioral model for reverberant modeling and an enhanced 3D experience.

The specific features include:

- Enhanced reverb effects added to the distance model to increase "realism"
- General Environment presets with global reverb and common material types
- Occlusion and Obstruction models for more complex interactive rendering

The IA-SIG interactive 3D audio guidelines are intended to influence how 3D technology is applied to PC systems (and interactive entertainment products in general) so that there is less variation among products and consistent use of terminology and technology. The desired result is consistent performance and less confusion. Besides representing industry consensus on the topic to the development community, the IA-SIG attempts to assist OS developers and reviewers (test labs, magazines, etc.) to understand and evaluate the differences between various approaches to 3D audio imaging, so as to best report their findings to consumers.

The initial documentation (which will be discussed at the 3DWG meeting at GDC) uses a DirectSound3D property set interface. This is because of the widespread nature of DirectSound3D, and as a practical matter for the members who were involved in specification creation. Once the final design specifications are complete, the group will work to generalize the specification into a platform-agnostic document.

The 3D working group expects to include a high-level API which can be used by ISV's and IHV's as a basis for implementation of the specification. There will also be proposed an appropriate acronym to refer to the functionality embodied by the Level 2 specification.

The level 2 guidelines are the combined efforts of many of the major players in the 3d audio hardware and software communities. However, special acknowledgment must go out to Creative Labs for their assistance in providing content from their Environmental Audio Extensions (EAX) to Direct Sound, and for contributing much of the text in the specification. The IASIG process serves as a model for cooperation by market competitors, working together to help raise the quality of interactive audio.

Submitted by Brian Schmidt <bschmidt@microsoft.com>
3DWG Steering Committee Representative

AUDIO ADVISORY WORKING GROUP (AAWG)

Chairman: Scott McNeese, VLSI <scott.mcneese@tempe.vlsi.com>

Overview

The Audio Advisory Working Group is focused on marketing efforts which will promote the expansion of the PC audio market.

Current Status

In 1998 the Audio Advisory Working Group (AAWG) focused on identifying tools and events to communicate and demonstrate the value of better audio. Activities included meeting with a game developer and exploring how a real application could be used as a demonstration platform and meeting with industry show management to investigate potential forums. The group's actions helped spawn two upcoming industry sessions designed to promote better audio. The first is Tom White's session at GameXecutive (GDC) titled "Profiting from Investment in Audio", and the second is an industry panel at E3 titled "Sound Strategy".

Future plans will be discussed at the next working group meeting in April.

<Ed: the difficulties the AAWG are encountering come as a disappointment, but not necessarily a surprise. In the crusade to better audio in games and other titles, a marketing strategy must be to convince project leaders and managers of the viability of excellent sound, which in itself is a very subjective thing. 3d audio does not identify itself as characteristically as 3d graphics. If the cause of the AAWG is to be furthered, a powerfully organized and well planned effort must be made on all fronts: the press, top officials, all the way to lead designers.

The question is: does the IA-SIG have the resources to accomplish this even if a plan is proposed? This could, and should in my opinion be a good question to ask at GDC. We must also remind ourselves that audio most certainly is getting more attention all the time, so the important thing to remember is that audio isn't being "ignored", merely given less importance than it deserves.>

INTERACTIVE COMPOSITION WORKING GROUP (ICWG)

Acting Chairman: Mark Miller, Harmonix Music <mark@harmonixmusic.com>

Overview

The Interactive Composition Working Group (ICWG) was formed in 1995 to explore and discuss issues relating to the creation and use of a family of technologies it refers to as Adaptive Audio Systems (AAS).

The ICWG defines Adaptive Audio as audio that is delivered via a system that allows for direct or indirect control of the data and/or data stream. A quote from ICWG member Thomas Dolby Robertson, Chief Beatnik of Headspace best describes the benefits of such technology:

"Adaptive audio systems provide a heightened user experience through a dynamic audio soundtrack which adapts to a variety of emotional and dramatic states resulting, perhaps, from choices the user makes." - Thomas Dolby Roberston

Update and Current Status

Two years ago, the work of the Interactive Composition Working Group was put on hold. At that time, the group was rapidly closing in on the definition of a generalized platform <and possible API> for the creation and playback of interactive audio. The group had arrived at a loose high level description of how such a system, called an 'Adaptive Audio System' would be structured. The next step would have been to begin prototyping. At this time, it became known the Microsoft was developing just such a system, now known as DirectMusic.

Rather than continue in parallel and invest in the specification of a competing API, it was decided to pause and let DirectMusic come to completion before proceeding. One of the technical leads had been an active participant in the ICWG unto that point and it was felt that the groups concerns and agenda had been successfully communicated to the DM team. Upon release, it was clear that this was true. Many of the terms and ideas discussed in the ICWG proceeding had, in fact, found their way into the system. This is, in fact, a great example of one of the IASIG's primary operating goals: 'to improve the performance of interactive applications by influencing hardware and software design'.

Now that DM has shipped, the group is being called to order once again. The main goals of the group going forward will be:

- To continue to foster the development and improvement of Adaptive Audio Systems across all platforms
- To advise developers of such systems of the needs and requirements of the users.
- To consider and document the likely future of such systems and their impact on the art of composition.

While I am currently serving as interim chairman of this group, I am actively looking for a talented a qualified individual to succeed me as my replacement.

Submitted by: Mark Miller

MULTI FORMAT WORKING GROUP (MFWG)

Chairman: Michael Land, LucasArts

Overview

The Multi-Format Workgroup (MFWG) was formed in the fall of '98 to study and define an important part of the interactive sound puzzle. Based on a brief but significant effort that took place at Project BBQ '98, the MFWG is charged with specifying which audio formats will be accepted as input by the computer's audio subsystem, and how they will be converted and mixed to whatever output format is required by the user. This is the place where things like 3D audio, quad, 5.1, home theatre, and various other technologies all meet in one big chart.

Current Status

Of course the real work (and the real value) is in the details. We've already discussed the basic input and output list, and we're about to start going over how the individual channels of each input format should be mapped to those of each output format. This is the real heart of the matter, the channel mapping for every input/output combination. When we're done in a couple of months, we'll have a blueprint for the next generation of interactive mixing engines, designed so that no matter what format of audio an author chooses to deliver, the audio subsystem will play it optimally on every system out there.

Submitted by: Michael Land

PLATFORM DEVELOPMENT WORKING GROUP (PDWG)

Chairman: Gary Johnson, Texas Instruments <g-johnson@ti.com>

Overview

The PDWG was proposed at the "Fat Man's BBQ" during October of 1996. At the BBQ, a prestigious group of content providers and computer audio professionals were asked "What do you want in PC audio in five years?"

The consensus was:

- Dedicated hardware acceleration for audio
- 32 channel digital mixing, with effects on each channel
- 6 channel mastering with multi-channel output or stereo output
- Multi-user interface
- Synthesizer - not FM, but Wave Table, DLS.
- Dedicated CODEC - this would be in addition to the AC97CODEC specification

A number of the hardware developers who attended the event continued the discussion of the matter on site and concluded that the components to the solution were actually available today, though not currently assembled into any one PC product or system. The PDWG was formed to further define a specification for this ideal audio platform, and to insure that a compatible, scaleable and expandable platform was created for the industry. The goal of the PDWG was defined as "To design and build a high quality audio platform(s) that content developers can create on, unrestrained by hardware".

To accomplish this goal, the group set the following agenda:

- (1) Define and design the platform, get further input from content developers.
- (2) Define and design a plan to evangelize the platform. This would include the creation of a logo, testing procedures, and methodology for accomplishing all of the above.
- (3) Evangelize the platform

The PDWG team spent the next year or so discussing specific means to accomplish the goals. As these discussions progressed, PDWG members took the ideas developed in the WG back to their respective companies. As most PDWG members worked on teams that designed audio platforms, this natural dissemination of WG information contributed to the independent manufacture of solutions that satisfied the WG requirements. As a result, the PDWG as a group was not required to actually develop a reference solution (as originally suggested) because the specified hardware began to appear in the marketplace, designed manufactured and distributed independently by various IA-SIG member companies.. Examples of such products are the SoundBlaster Live!, the Aureal Vortex 2, and the Yamaha YFM724.

In conclusion, the PDWG has accomplished what it set out to do at the '96 BBQ in a profitable and commercial fashion. Thus while no tangible reports or official guidelines resulted, the work of the WG led to industry progress in its area of focus and thus can be considered to have achieved its goals successfully. That said, the WG does not plan further activities at this time and the Chairman recommends to the Steering Committee that the WG be closed.

Submitted by: Gary Johnson

Section IV. COLUMNS

This section contains features and columns on varied topics of IA-SIG interest.

=====

Interactive Composition Column 1.1

A Direct Music Primer
by Alexander Brandon and Mark Miller

Much has been said in the past few months about DirectMusic and, thereby, about interactive music. In this first column, we will provide a basic outline of DirectMusic and then lay out some real world and theoretical context in which to examine it. In future, we will delve down into the details of actually using the system and others like it to create interactive sound tracks.

Introduction

DirectMusic is a system native to the DirectX 6.1 SDK which adds specialized music playback and synthesis to the next generation of the Microsoft "Windows" operating system. DirectMusic Producer is the editing environment in which the composer can write music and create different "Styles" in which the music can be changed interactively. Rather than simply play back preset MIDI or Digital Audio data, the DM system works with a set of parameters defining music content generation. In this context, 'generation' refers to the fact that the actual music that is played back is assembled by the playback engine, from a combination of these parameters, music data created by the composer and input about the state of the application (a game, for example) just before it is needed for playback. This makes the music more "interactive" and responsive.

The system also improves PC sound playback quality by standardizing a synthesizer in software rather than relying on the varied existing hardware specifications and solutions, (although hardware specific channel playback is supported.) The system also adds variety and power by supporting Downloadable Sounds (DLS1). Beyond what is shipped on the disk, the most interesting and perhaps most unique

attribute of DM is its open ended architecture. This architecture allows developers to create applications such as sequencers, MIDI effects processors and custom sound design tools to further simplify and or enhance the composition and playback process.

So you can see that DirectMusic is not the end of PC composition, it is merely a more solid foundation for music playback and construction methods going forward.

Design Commentary

The system is quite an interesting and imaginative leap for Microsoft, but not one that comes as a surprise. For the last several years Microsoft has expanded its operations in nearly every field of computing, however its MS-DOS and Windows offerings have not been revolutionary, rather they have emulated and built upon them existing trends, like the MAC OS, (for example). Apart from the DirectX SDK, DirectMusic is the first such example we have seen from MS of a completely new approach to generating content. It will be met with a great deal of enthusiasm as well as a great deal of criticism. For the purposes of this report, the product will be simply be reviewed and thoroughly examined on it's own merits.

The only Microsoft specific consideration to be made here is the fact that DM will have a much wider market distribution than most Interactive Music products, being that it is included free with DirectX 6.1 SDK. Considering the number of Windows machines in out there today, and their use by nearly every developer creating software, this fact must be taken into consideration. Before we can really judge DM as a product, we must first present an approach to interactive music in games, a rapidly expanding and evolving field, and compare it's requirements to the capabilities of DirectMusic as closely and thoroughly as possible, being precise as to details.

At first glance, DirectMusic and DirectMusic Producer comprise an excellent set of tools for developers. That being said, we must ask the question "how excellent?" Considering other, established tools such as Miles, Beatnik, AMStudio, and all of the proprietary ways of approaching interactive music will give us a platform for comparison. (details on systems such as LucasArts' "Imuse" and EA's Adaptive Audio engines will follow in future updates, for this purpose). But before we do even that, we need to answer "The Big Question":

Has Microsoft left the station too late?

Most, if not all industry veterans agree that if DirectMusic had been presented five years ago it would have been the best method for game music development and been rapidly adopted. Today, however, new methods of real-time music streaming from CD / DVD are being implemented that make DirectMusic seem obsolete, so is this approach still valid? After serious consideration, the answer must be, yes. Real-time multiple track music streaming can be a great way to do game sound tracks, but it certainly isn't the only way, or even necessarily the best way. One can stream music, give a game a movie-like feel with this technique and come off with a superlative soundtrack, but the music can't be varied responsively each time the user plays. This static quality often detracts form the player's satisfaction. Streamed digital audio is also large <storage wise> and bandwidth intensive, which can severely limit the scope of its application. This would lead one to ask:

How important .is. actual sound playback quality in the overall picture?

This is a very important question indeed. Today, it could be argued that MIDI controlled playback on soundcard technology still can hardly hold a candle to professionally recorded audio. If this is true, than one could not fault a composer for wanting the highest standards possible and wanting to go pure DA. But the game industry is still a leap from the movie and music industries in terms of playback system technology. Compare, for example, a typical pair of \$20 PC speakers to a home or actual theater sound system. Given this and the new capabilities presented by DirectMusic the gap may will be narrowed even

further. With its support of custom DLS based samples instead of stock GM presets, DM's interactive MIDI playback should finally provide a listening experience that can approach the quality pre-recorded music. This is, in fact, still an area where an under-funded practitioner can begin with the simplest of tools and create acceptably impressive work. That said, there are still some major discrepancies. What about, for example, real time DSP effects? dynamics? EQ? Stereo enhancement? Digital 5.1 surround? The road to outstanding sound quality will not be a short one, but DirectMusic has at least put its foot down in the name of interactive music where others have been loathe to tread.

Lastly, this streaming DA vs. interactive MIDI plus DLS question is really comparing apples to oranges. The only problem in this comparison is that people are used to apples, and while the orange may be a better tasting fruit, people won't know until they actually have a taste. Streamed music in the form of pre-composed performances has been the staple of listening and music in general throughout human history. Attempts at randomization (Cage, for example, and Stockhausen) may have been recognized as leaps forward, but have failed to become accepted in popular culture. So one might say that DM, the untested fruit, will really need to prove itself powerful and useful in the hands of those composers who embrace it, or be left behind.

Even so, many game developers have already begun planning for adding DirectMusic to their games. Other game audio system projects, both established and in the works, are planning to apply similar methods of interactivity as those used by DM to their offerings. So while DirectMusic may not have the highest standard of sound quality compared to CD / DVD streaming, there seem to be compelling advantages to its approach that are gaining a large foothold in the developer marketplace. Perhaps, then, the appropriate behavior of music in an interactive context may, in fact, be equally or more important than the bottom line sound quality.

Interactive Audio in Games - An Approach

In games, when we look at how music is to be played, we see several methods. The most conventional is a single looped piece for each "level", (or whatever term is described for a chunk of playtime that is defined by the developer), being played constantly over and over again. This has worked quite well for many years, but the fact that the public seems satisfied with it does not mean it is the way to use music most effectively. Music, being so abstract and nebulous, does not lend itself to straight forward paths of evolution in perceived quality the way video technology does in the improvement of image quality. 3d sound, for instance, does not have quite the immediate and awe inspiring effect of 3d video.. (genuine 3d video, holographic, etc..).

What we can now do is look at different ways to play audio. Gathering information from years of subject discussion, we can see that random music generation in its purest form is just as nonsensical as random visual generation. People need a foothold in reality, things must have consistent shape and color definition to be held as discernible objects. So let us toss out completely random music generation for now. It may be used in the future, but we do not yet know a method to make this practical and useful in today's context. A more fruitful approach would be to begin with music that has recognizable elements of shape and form. This music should then be adapted during playback to fit changes in user action and location. This method has already been used in a variety of contexts, with some notable success. This is what we will pursue for now.

Thematic Development

To begin this discussion, we must talk about thematic development. Thematic development is something that has been vital to successful music scores, most recently, game scores, for many hundreds of years. We will take a moment to explain it so that we can fully understand its importance. (NOTE: This is a theory and not a fact about music. Music holds few, if any, absolutes, as any art form does, so we are certainly not expecting everyone to agree with the following analysis.)

To identify with something readily, there must be something constant and unchanging in it. There must be at least one identifiable characteristic that remains in the memory from encounter to encounter. Objects currently existing in reality such as trees and birds present little problem with this kind of recognition. Used in virtual environments like games, such representations and variations of these real world shapes and objects give the player a start on a comfortable and recognizable set of surroundings.

Once that is done, the unimaginable and non-repetitive can be ventured into. This is, however, very difficult to do well and is based on relativistic principles that we won't go into here. This is not just something that is done in games, but in movies, books, and nearly every leap of the imagination presented to a wide audience. Science Fiction, for instance, presents bold and incredible concepts, but by the rules it is given must base those concepts on things already discovered. This may sound limiting, but in fact it is essential for the readers' or viewers' enjoyment.

In a novel, the author begins by using something already understood... English, for example, to define things the reader can recognize... the language we know speaks of planets, stars, humans, human behavior, etc.. And then the author begins to unravel this comfortable blanket of common knowledge around the reader to expose them to things they have never before considered... Doing so leaves an unforgettable effect if done with care and precision.. The best writers of fiction, fantasy, and science fiction have proven the success of this method for many years.

The same set of rules can be applied to interactive game music in a loose way. By giving the environments and characters in a game themes, the player can grasp their identity, and once the themes are established, variation can be introduced. Easy examples are movie soundtracks, and symphonic classical and romantic orchestral music, from Musorgsky's "Pictures at an Exhibition" to John Williams' "Jaws".

Repetition in contrast to theme and variation

In an article by David Yackley and a summary of DirectMusic elsewhere in Microsoft's documentation, the author stated that "repetition is boring". In one sense, this is correct. In another, it is not. The authors aren't guilty of stating that DirectMusic should vary everything all the time, but they don't necessarily paint the whole picture. Developing and establishing a theme before introducing variation, while not absolutely necessary, adds a great deal. It familiarizes the player with the game in a more intimate way than constant, pervasive variation can.

On the other hand, repetition in its purest sense will bore the player eventually. The answer to this is DirectMusic's ChordMaps, Styles, and Templates. When used properly, the composer can take a theme for a character or an environment, specify the harmonies they would like to see used, add rhythm if necessary, specify the variation, and let the system create the desired effect. In this way, the soundtrack can STILL have the themes that a composer / producer wants for the characters, and yet programmatically introduce meaningful variations.. Examples of this will be forthcoming.

In conclusion, the idea of thematic development, of buildup and resolution, are very important to the successful realization of real-time, interactive music. Currently, DirectMusic tackles this in a more complicated and comprehensive way than ever before.

In the next installment of this column, we will begin to go into the details of how this is actually done, in DM and other Adaptive Audio systems.

<End part one>

Section V. INDUSTRY CORNER

=====

The bold industry is what drives development, and development drives the products. Paying close attention to the industry has been a very important duty for the IA-SIG and we will begin to document that watchful eye in this section of the newsletter. Read on and find out just what the blazing people out there are doing with your tools, both in software and hardware.

"Engine Roundup"

by Alexander Brandon

With the advancements in PC audio made over the last five, even last two years, it is becoming clear that we are no longer dealing with merely FM sound and wavetable cards. We have massively powerful 3d surround engines that add realism and dynamics to our games and apps. We also have many different kinds of music engines that support many different kinds of playback, from Redbook to multichannel WAV to MP3 to MOD to DLS / MIDI, the well informed audio developer is now awash in a veritable sea of possibilities.

But the operative phrase here is "well informed." This report is intended to help clear up some of the confusion as to just what the techniques for generating state of the art audio content are and what options are available to the developer. While it should be especially helpful for the layman, hopefully there is also information for the seasoned composer/programmer/engineer that will be of interest.

We will begin by listing engines that are intended for music playback. Some engines feature both music and sound effects playback routines, some engines feature various ways of doing one or the other, and some actually contain production applications. Integration of separate engines will be featured in a later issue along with specific examples of published games where each engine has been used.

(ed: all information was provided by the companies themselves or from public sources. The IASIG is not responsible for the accuracy of this information and does not verify nor endorse any of these products.)

STACCATO SDK

Staccato Systems, Inc.
Palo Alto, CA
(650) 853-7035
<http://www.staccatosys.com>
(Windows)

Overview: Commercial Software DevelopersKit (SDK) includes an algorithm engine for games and a license to ship for commercial apps.

Features: Staccato provides tools that empower sound designers to develop audio synthesis and processing algorithms, a flexible host-based synthesis engine to run them, and content that shows the power and flexibility of the audio rendering engine. Staccato has a large amount of intellectual property to bring to the audio market, particularly in the area of "Physical Modeling" synthesis and other "next-generation" audio rendering techniques. The Staccato audio rendering engine, "SynthCore" also allows for other forms of synthesis such as wavetable/sample-based synthesis as well. Staccato brings some real advantages to the user experience and development process by:

- Providing maximum audio rendering versatility via multiple synthesis methods and modular architecture
- Allowing you to use your current sound libraries as well as develop new, more differentiated sounds
- Letting the game physics and environment drive the audio in a more interactive way
- A predictable user experience, since the solution is not audio hardware-specific

SynthCore(TM) is a synthesis engine that is Sordius-XG, XG-Lite, GM, Direct X, Down Loadable Sounds (DLS) and Down Loadable Algorithm (DLA) compatible.

MILES SOUND SYSTEM 5

Rad Game Tools, Inc.
Kirkland, WA
(425) 893-4300
<http://www.radgametools.com/miles.htm>
(Windows)

Overview: Perhaps the most used commercial audio engine on the market. Miles Sound System provides MIDI, XMIDI, Redbook, MP3, and DLS support including a software synth and integrated support for numerous compression schemes.

Features:

- Dolby Pro Logic compatible surround
- Intel based RSX 3D positional sound
- Mpeg layer 3 decompression support
- High level 3D Audio API
- New digital subsystem
- DLS level 1 support

DIRECTMUSIC

Microsoft Corporation
Redmond, WA
(425) 882-8080
Contact: Chanel Summers <chanel@microsoft.com>
(Windows)

Overview: DirectMusic (DM) is a brand new and powerful computer music development tool using DLS as its primary delivery method for sounds and MIDI as its control. The DirectMusic SDK is available free with DirectX 6.1 SDK, along with "DirectMusic Producer", the editor for DirectMusic. Information can be found on Microsoft's DirectX site, and in this issue's column on interactive composition.

Features:

- Uses DLS level 1
- Imports MIDI files
- Expands interactivity through use of Variations and Chord Maps
- Realtime reverb and chorus
- Unique composition environment found in DirectMusic Producer.

GALAXY SOUND ENGINE

Digital Dreams Multimedia
Contact: Carlo Vogelsang <carlo@ddreams.nl>
(Windows)

Overview: Galaxy is a state of the art music playback system featuring Mpeg layer 1, 2, and 3 support, MOD file support, Redbook, and MIDI support.

Features:

- All major format support

- Support for most soundcards, including Adlib Gold 1000 for DOS based applications, up to the latest MME compliant cards
- Support for EAX and A3D for sound effects playback

HEADSPACE AUDIO ENGINE (HAE)

Headspace, Inc.
 Contact: Steve Hales <steve@headspace.com>
 (Mac / Windows (playback only))

Overview: The Headspace Audio Engine is designed to effectively play music and sounds over low bandwidths such as the internet, and can also be used in other applications such as games. The engine contains a linear audio API, a file reader, a software based wavetable synth and sequencer, and a 64 input, 2 bus mixer. The engine will support files from numerous sources including Headspace's own .RMF file format also known as files created in their music editor "Beatnik".

Features:

- Supported files include MIDI, RMF (native Beatnik files), MOD, WAV, AIFF, and AU.
- Supported platforms include WebTV, JavaSoft JVM, Sun Solaris, BeOS, Netscape, Magic Cap, and others currently in alpha and beta stages (Director Extra, ActiveX, Win95 VxD)

PATHFINDER (Proprietary engine used by Electronic Arts only)

Contact Alistair Hirst for more information: <ahirst@ea.com>
 (Mac)

Need for Speed II and III featured interactive streamed music on both the PC and Playstation. By streaming, the music was able to maintain the production values of music produced in a professional studio, and leave more sound RAM for sound effects.

Interactivity was achieved by seeking around in a linear music file. The music was written in sections which both attempted to reflect the surroundings of the track where they would play, and in various levels of intensity which would reflect the performance of the person playing the game. Pathfinder is a Mac based proprietary tool used at EA Canada and EA Seattle to set up a tree to reflect how the different sections of the music were related, and to set up a "path" for the music to follow based on input from the game. It also allows the setting of triggers, which when called, will jump the pointer to a different part of the stream. Pathfinder's output is recognized by the in-house libraries for both platforms, making integration simple.

Section VI. Developers Corner

At last, a place where the developers can speak their minds. The interactive audio industry is bursting with men and women ready to both complain and cheer about various aspects of their work. Here is an outlet for them.

"Working on Myth"

by Marty O'Donnell, Total Audio (<http://www.TotalAudio.com>)

The interactive audio engine used in both Myth games was an extremely powerful and probably under utilized engine. Because the developers didn't want game runtime "music" we only used it for sound design. However, it had some incredible features that could have just as easily been applied to interactive

playback of music. As a matter of fact the "loss" screen in one place where I experimented with it's use. I'd love to expand on it's features because I think it could be applied to every audio engine out there. It has to do with concepts rather than technology and deals with recordings rather than MIDI. The future of robust interactive audio has to include this kind of solution as well as MIDI solutions. By the way, we still had 53 minutes of recorded music in Myth II Soulblighter.

"Top ten things a video game producer should know about creating great audio"

I recently requested input from you all on this topic for a paper that I was working on for GDC. I got some great quotes and thought that I would publish them here.

- Mark Miller

- "Just because it sounds like your favorite band doesn't mean it's the right thing for this project."
 - "No, you can't play the guitar solo."
 - "Listening to a lot of records isn't the same as knowing how to make one."
 - "Quality costs money. If you're doing it cheap you're either not getting quality or you're ripping someone off."
 - "Tell me the emotion you want to create, not how to write the drum part."
- Spencer Critchley, Director of Production, Silicon Gaming, Inc. <spencer@ubet.com>
-
- To quote our dear compatriot Brian Coburn: "People don't really like music, they just expect it..." That just might be near number one(!)
- David Javelosa, 31st Century Noise Design, <davidja@brandx.net>
-
- 1) MIDI is a control protocol, not a sound. ("What you hate is the sample set, not the MIDI.")
 - 2) You'll get better sound for the same money if you amortize audio department infrastructure across multiple titles rather than expensing against individual projects. ("We know the testers are logging bugs on SGIs, but we just paid for all those massages, so we can't afford that used PowerMac 8500 - and- the AudioMedia III. Can't you just do the game with SoundApp instead?")
 - 3) Avoid build processes and APIs that make the sound artist go through engineers to make sound-only changes.
 - 4) Just because your cool 3D world geometry engine looks great because it's physically accurate doesn't mean adding a geometrically accurate sound model is gonna sound right. ("But the box says it's - accurate-! It -can't-sound bad!!!")
 - 5) Don't put too many eggs in the 3D audio basket, game-design-wise. ("Y'know, with only two speakers, only the .helicopters. can go behind you...")
 - 6) Latency .really. matters; a button click sfx heard 200 mS after you click the mouse isn't a button click any more. ("200 mS isn't a delay, it's a postponement!!")
 - 7) When all games are using the same sample set you are, there's a limit to how much better you can make your game sound. ("Budget realistically for custom samples!!")
 - 8) Dialog recording sessions are like motion capture sessions: turning an actor's output into usable elements can take a -lot- of cleanup work. ("Budget realistically for dialog editing!!")
 - 9) When the programmer says that what the sound guy wants is impossible, the programmer is frequently, er, lying. ("How can you tell if the programmer's lying to the sound guy? His lips are moving!!!")
 - 10) Allow calendar time in the schedule to review the sounds in context (with real players) for a while, and then get one or two rounds of revisions in. A sound design concept takes calendar time to ripen, and it can take some hours to smooth rough edges in the first implementations.
- Chris Grigg <chrisg@sirius.com>

COMING IN THE NEXT ISSUE (June 1999)

- Feature on interactive 3d positional sound engines including EAX, Qsound, A3D, Focal Point, and more
- Interactive Composition Column part 2
- Working Group updates
- A lot more developer comments and industry updates
