

Session A-3
Preprint 4023

Huron - A Digital Audio Convolution Workstation.

Presenters: David McGrath, Andrew Reilly
Lake DSP, Maroubra, NSW

Presented at
5th Australian
Regional Convention
April 26 - 28, 1995
Sydney



AES

This preprint has been reproduced from the author's advance manuscript, without editing, corrections or consideration by the Review Board. The AES takes no responsibility for the contents.

Additional preprints may be obtained by sending request and remittance to the Audio Engineering Society, 60 East 42nd Street, New York 10165 USA.

All rights reserved. Reproduction of this preprint, or any portion thereof, is not permitted without direct permission from the Journal of the Audio Engineering Society.

AN AUDIO ENGINEERING SOCIETY PREPRINT

Huron – A Digital Audio Convolution Workstation.

David McGrath
Lake DSP Pty. Ltd.
4/166 Maroubra Rd.
Maroubra 2035

dmcgrath@extro.ucc.su.oz.au

ABSTRACT

After several years of development of Audio DSP processors, the demand for complex DSP systems with multiple channels of I/O has led Lake DSP to create a Digital Audio Convolution Workstation. This system, known as Huron, is intended for use in research areas of audio, where a large array of parallel DSP processors can be used to process multiple channels of real-time audio. The Huron system is built around a powerful host processor with an integral audio bus that allows 256 channels of audio data to be passed between DSP processors and I/O modules. I/O options include A/D, D/A and AES/EBU digital I/O. Huron is also the first platform available that uses the zero-latency convolution algorithms developed by Lake for real-time audio processing with ultra-long FIR filters, up to 262144-taps on multiple channels. The Huron system is also capable of rapidly updating long FIR filters to allow real-time animation of complex effects. The main applications addressed by Huron include 3-D virtual reality audio, accurate reverberation, equalisation and multi-channel mixing for studio or sound-reinforcement systems.

0. Introduction

This paper presents an overview of the features and applications of the Lake DSP “Huron” digital audio convolution workstation. In section 1 we see how Huron meets the needs of real-time audio work. Section 2 describes the IO system, and the various options for attaching input and output devices to Huron. Section 3 discusses the concept of convolution, and sections 4 to 7 focus on four key application areas targeted by Huron.

1. Requirements of real-time audio

In this section, we examine the requirements of a real-time audio workstation. We see that Huron meets all of these requirements for a wide range of audio applications.

1.1 Low overhead input/output

Signal processing algorithms require timely access to inputs and outputs. Some applications access a very large number of inputs and outputs. In order to maximise the effectiveness of the DSP processors in a real-time audio system, it is desirable to minimise the number of cycles required by the DSP to access successive input values, and to output

results, because every cycle spent transmitting or receiving data is a cycle that can not be used to calculate the result.

The Huron bus architecture provides 255 channels of digital audio data bandwidth, for input and output signals. The processors access this bus through a pipelined address mechanism that allows input and output operations to occur in a single DSP instruction cycle, and because of the pipelining, successive input and output operations can occur back to back, cycle by cycle. This is a significant improvement over some other systems, where the address for each input value must be determined and used by the processor in order to access that value.

1.2 Easy configurability

The Huron bus address generation mentioned in the preceding chapter has other significant benefits at the system level. Since the channel number of input and output signals are generated by the Huron bus interface, and can be modified while programs are running, signal processing programs can be written as functional, re-useable blocks, without concern for their run-time configuration. For example, a DSP program that performs a mixing function does not require any modification for operation with inputs 2,3,4,5 instead of inputs 10,11,12,13 — these input and output assignments can be made and re-made at run time.

1.3 Sufficient computation resources

Some signal processing operations are very simple, and their only operational requirement is a real-time processor of some sort, and perhaps a little memory. Into this category one might put mixers, level meters, fixed delay units, and other traditional musical effects. For these applications the ability to handle multiple channels of I/O is paramount. There are other applications, such as spectrum analysis and convolution, that are extremely computationally expensive. In fact it is only by being able to provide such a high concentration of DSP power as that available in Huron that these applications are possible at all.

The Huron system, as a result of its flexible digital audio bus structure, can be provided with as much processing power as a particular application requires. The minimum configuration is one processor card, which consists of four Motorola 56002 processors, each of which has access to 128k words (384k bytes) of high-speed SRAM and 1M word (3M bytes) of DRAM. As well as access to the 256 channels of audio data on the Huron backplane, there are an additional 512 channels available locally on each card, for inter-processor communication within the four processors on a card. This one card is sufficient to operate simple mixing functions on all audio channels, or provide long, zero latency convolution at 48kHz, on one or two channels.

1.4 Low overhead communication between parallel processes

The processors in a Huron system can communicate with each other simply by sending and receiving values on Huron bus channels that are not occupied by A/D, D/A or digital I/O modules. In most systems, this represents the bulk of the backplane bandwidth. As mentioned in the preceding sections, processes that operate on DSPs on individual cards have three times as many channels to use, and (just as with signal inputs and outputs) this data communication costs only a single processor cycle.

2. Signal input and output modules

The Huron IO sub-system is based on an IO-module backplane card that can carry up to eight stereo modules. Each module can be either analog input or output or digital input or output. The details of each module type are discussed below.

2.1 Analog input modules

The analog input modules provide very low noise, 16-bit conversion of analog inputs to digital signals, which are driven onto the Huron backplane bus. For maximum flexibility, each analog input is also provided with a digitally controlled, analog volume control and variable gain pre-amplifier. These allow signals ranging from professional line level equipment to dynamic microphones and electric guitars to be plugged straight into the XLR connectors provided on the Huron patch panels.

2.2 Analog output modules

The analog output modules provide high quality, 18-bit conversion from digital data on the Huron bus to analog signals. These outputs are also provided with a digitally controlled analog volume control, so that the output level can be accurately matched to the input level requirements of the amplifier, headphones, or recording equipment that is plugged into the Huron patch panels.

2.3 Digital input modules

The digital input modules accept AES/EBU standard digital signals at 48kHz, 44.1kHz and 32kHz. The system clocks can be derived from one of the digital inputs if required. The digital inputs provide a mechanism for direct input of 24-bit audio data, the maximum resolution available with the Huron processor cards.

2.4 Digital output modules

The digital output modules also drive AES/EBU standard digital signals at the system clock rate.

3. Applications of convolution

Convolution is a mathematical term that describes the operation of generalised filters in abstract linear systems. Until the development of the Lake Huron Convolution workstation, convolution had only minimal impact in practical signal processing applications, because it is inherently computationally expensive. Some small filters, such as the anti-aliasing filters in analog-to-digital converters are implemented with convolution, by using specially designed hardware. These filters have around 32 coefficients, and so can only perform this simple, low-pass filtering function. In the following sections, the relationship between convolution and filters of all sorts will be explained, and then the applications that become possible when very long convolution is available will be examined.

3.1 Linear systems and filters

To a close approximation, most real-world systems behave linearly. That is, for every input, there is a proportional output. When real world systems have some memory or

state, where observable “outputs” can occur after the associated “input”, and where this output happens whenever the same input is used, then the system is called a “linear, shift-invariant and causal” system, and it becomes very easy to describe mathematically. In such systems, it is possible to describe the relationship between inputs and outputs in terms of a “transfer function” (in the frequency domain) or an impulse response (for time-domain signals). Mathematically, we write convolution as:

$$y(t) = x(t) * h(t), \quad (1)$$

where $y(t)$ is the signal at the output, $x(t)$ is the input signal, and $h(t)$ is the “impulse response” of the system.

3.2 Impulse responses

The impulse response is a time domain version of the transfer function (which is usually expressed as a function of frequency). In a digital implementation of this system, with discrete-time signals and a finite-length impulse response consisting of N samples, this expression can be re-written as:

$$y(n) = \sum_{i=0}^{N-1} x(n-i) h(i). \quad (2)$$

That is, for every sample of the input signal, multiply it and each of the preceding $N-1$ samples by corresponding coefficients from the impulse response, and add them all together. The result will be the output of the system for that time sample.

Impulse responses can be measured or synthesised. Where the linear systems under consideration are the acoustic responses of rooms, then the impulse responses can be very long. For large halls or cathedrals, responses of several seconds are not unusual. In a digitally synthesised filter, this duration corresponds to a very large number of coefficients $h(n)$.

3.3 Convolution is the appropriate tool

In spite of the computational complexity, convolution is the only mechanism by which effects such as accurate reverberation, or arbitrarily complicated equalisations can be performed in real time. Simpler structures, such as IIR (infinite impulse response) filters have been the mainstay of equalisers, because they can be constructed to operate in real time. Where long, complicated impulses have been required, in architectural acoustics, for example, then the only alternative has been to compute the convolution off-line as a batch process. With the advent of zero-latency long convolution, this is no longer the case. Direct use of convolution is the most desirable (perhaps only) method.

4 Applications: Virtual reality audio

In the emerging field of virtual reality, where the goal is to create a believable synthesis of sensory inputs for the human operator, the complexities of the auditory inputs are popularly overshadowed by the exciting animated graphics. However the acoustic simulation of the virtual world is just as complicated and requires similar levels of computational

power. Not only is it necessary to make sound sources appear in different locations in space, but the echoes and reverberations of those sounds must be realistic, and the whole experience must twist and turn in response to the user's head movements.

With the Huron, we have developed systems that solve most of these problems, and have demonstrated moving sound sources, with head-movement tracking and realistic room reverberation.

5. Applications: Accurate reverberation

As described in the next paper (Reilly and McGrath, 1995) it is possible to use long convolution to exactly model the acoustic path between a sound source and a listener's ear drums. That is, if a recording is made in an anechoic environment and then convolved with the impulse response of a particular hall or room, then the result will sound exactly as though the listener was hearing the original performance in that hall or room. This is a very valuable technique for architectural acoustics, where one might like to compare the sound of two halls, or compare a hall with the addition of some new baffles, before they are physically built. The first of these applications requires the measurement of the impulse responses of the two halls in question. The second requires accurate synthesis of the hall and the baffles.

5.1 Ability to measure room reverberation

The technique for measuring room reverberation involves exciting the room with a white noise source, recording the result using a desired microphone or listener position set up, and then de-convolving the white-noise, to leave the impulse response of the room. This process is described more fully in the next paper. This ability to measure the impulse response of a room is a standard feature of the software sold with the Huron system.

5.2 Ability to synthesise room reverberations

A more complicated problem is that of impulse response synthesis. This is typically approached in the same way that image ray-tracing is computed. The complication with acoustics is that it is not hard to have object features that are of similar sizes to the wavelengths under consideration, and this causes diffusion and refraction, which are very difficult to model. Several sources for room impulse response synthesis are listed in the references.

6. Applications: Equalisation

Equalisation is a term that is often applied to a generalised filtering process, where that process is intended to remove an unwanted component of a sound source, or to cut or boost a range of frequencies.

6.1 Graphic and parametric equalisation

The traditional devices used for equalisation are "graphic equalisers" and "parametric equalisers". Graphic equalisers consist of a collection (often 30) of band-pass filters with variable gains. These can be used to cut or boost frequency components in octave or third octave bands. They are simple to use, but do not offer precise control of the filtering process, and often introduce significant extra noise. Parametric equalisers consist of a

collection (often only three to five) of band-pass filters that can cut or boost a group of frequencies, but in this case the center frequencies and bandwidths are not fixed. Much greater control is therefore possible with a parametric equaliser, but they can be difficult to use, as one filter band can influence the performance of another. Also, since the number of bands is often limited, the number of features of the sound that can be corrected is correspondingly limited. Another problem with both types of equaliser is that they introduce uneven group delay into the signal.

6.2 Equalisation by convolution

Since both of these types of equalisers are simply filters, they can be implemented by convolution. All that is necessary is to take the user's desired frequency response curve and synthesise an impulse response that produces that result. The Huron EQ software does exactly this. It allows users to specify equalisation curves in terms of bandpass, highpass or lowpass filters and specify their bandwidths and gains, just as they would with a parametric EQ. There are a number of significant advantages to the Huron method, though:

- The whole equalisation response curve is continuously displayed on the screen, so interactions between filters are clearly visible.
- The equalisation curve can be arbitrarily complex—there is no limit to the number of “parametric” components that can be added.
- The user now has independent control of the phase response of the filter. It is possible to use a traditional, minimum phase response, or the digitally perfect “linear phase” response (that incurs a significant fixed group delay), or a user-selected group delay that combines the best characteristics of each method.
- The filter introduces no noise, no matter how complex the response. The convolution filtering is implemented with high precision fixed point arithmetic which keeps computation noise well below the level of the audio data.

7. Applications: Multi-channel mixing

Mixing was mentioned in the introduction as another application of the Huron platform. This application relies on the high data bandwidth and ability to configure the system with many input and output channels, rather than its ability to perform convolution.

When there are many channels, all of which may require attenuation and metering, and mixing onto output channels, it is very important that both the hardware and the software have support for this. At the hardware level, Huron has the low-overhead access to the channel data and the on-the-fly re-configurability required. This was discussed in section 1. Huron also has support for multiple channels at the software level too. The software allows for naming of each channel, and storage of channel configurations on disk. The “Virtual Rack” user interface has a flexible and familiar appearance—it is just like plugging equipment and cables into an equipment rack.

8. Conclusion

In this paper we have described the design and capabilities of the Huron digital audio convolution workstation. We have shown how it meets the needs of audio professionals in a wide range of applications, and for a wide range of problem sizes within those appli-

cation areas. Convolution has been shown to be the best tool for a wide range of filtering and simulation tasks, and that Huron can perform this convolution with zero latency, using long impulse response lengths, simultaneously on multiple channels of audio data.

References

Acoustic Design Ahnert, *ADA EASE*, Gustav-Myer-Allee, 25, D-13355 Berlin.

Acoustics Laboratory, Institute of Electronic Systems, Aalborg University, Denmark.

BOSE Corporation, *Auditioner*, 100 Mountain Rd, Framingham MA 01701, USA

Computer Aided Theatre Technique: attn: B-I Dalenbäck, (1994) *CATT-Acoustic version 5.0: Room Acoustic Prediction and Desktop Auralization*, Mariagatan 16A, S-414 71 Gothenburg, Sweden.

McGrath, D., (1994) "A new approach to digital audio equalization", in Proceedings of the 97th AES Convention, San Diego.

Polhemus (1993), *3SPACE InsideTRAK*, One Hercules Drive, P.O. Box 560, Colchester VT 05446, USA.

Reilly, A. P., and McGrath D. S., (1995), "Convolution Processing for Realistic Reverberation." in Proceedings of the 98th AES Convention, Paris.

Timeware Corporation, *Realwave*, Nakanobu Bldg, 2-4-10, Higashi-nakanobu, Shinigawaku, Tokyo 142, Japan